Canadian Journal of Philosophy (2026), 56: e21, 1–19
doi:10.1017/can.2026.10113

SYMPOSIUM ARTICLE

How Climate Change Confounds Conventional Normative Theories

Joseph Heath

University of Toronto, Canada
Email: joseph.heath@utoronto.ca

Abstract
Normative theories often fail when applied to future generations, creating counterintuitive results in climate change policy. This paper explores these challenges to explain why many sophisticated arguments initially seem implausible. It divides normative ethics into consequentialist and deontological approaches, showing how both schools struggle to provide a coherent framework for intergenerational justice and climate change.

Keywords: climate change ethics; future generations; sustainability; demandingness; discounting

One of the most controversial features of the 1996 Kyoto Protocol, which sought to institutionalize a set of "common but differentiated responsibilities" to mitigate global climate change, was that developing countries, including China, were exempted from having to make any reductions in emissions. Since that time, Chinese CO_2 emissions have increased from just under 4 million gT per year to over 12 million, making that country by far the largest source of greenhouse gas emissions in the world (Crippa et al., 2024). (By comparison, total emissions in the United States have remained relatively flat for the past decade at between 5 and 6 gT.) The argument of the government of China, back in 1996, was not that the country should be free from any obligation to reduce its emissions, but rather that it should be permitted to prioritize economic growth in the near term. In other words, it was an argument for *delay*. The Chinese strategy was to maximize economic growth first, using all available means to expand usable energy supply, then later begin transitioning to carbon neutral energy sources.[1]

Although the position of the government of China was met with considerable derision at the time, both by environmentalists and US politicians, in retrospect it is very far from obvious that it was wrong.[2] Average real per capita GDP in China, which was $709 in 1996, had risen to $12,720 by

[1]As Darryl Moellendorf (2014), 17) observes, "energy poverty" is a major barrier to development in the world today. He offers the following summary of the basic tradeoff involved in climate change policy: "If mitigating climate change involves slowing the process by which billions of people will climb out of poverty, perhaps people in the future are better off with less mitigation. Overcoming energy poverty today will improve the lives of people in the future. Putting ourselves in their shoes, we judge there to be good reason not to want to retard or delay the long-term process of human development by increasing energy costs, a process that will yield capital investments from which they will benefit" (Moellendorf, 2014, 18).

[2]Exclusion of China and India from the list of Annex 1 countries was the official reason given by U.S. President George W. Bush for the U.S. refusal to ratify the treaty (Berger, 2001).

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Canadian Journal of Philosophy, Inc. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial licence (http://creativecommons.org/licenses/by-nc/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use.

2022.[3] This has been accompanied by annual wage growth that has averaged 10% over the past two decades, resulting in vast segments of the population experiencing a radically improved standard of living (World Bank, 2022). Infant mortality in China, for example, has declined from 37 per 1000 live births in 1996 to less than 6 in 2020, resulting in approximately 375,000 fewer babies now dying per year.[4] (Maternal mortality has also declined by 80%, resulting in more than 7,000 fewer women dying in childbirth each year [Chen et al., 2021].) Malnutrition went from being a widespread domestic problem to having essentially been eliminated. Furthermore, the expressed commitment to making a later transition to carbon neutral energy production was not mere rhetoric. Since 1996 China has expanded its number of nuclear power generation facilities from 1 to 58, with another 30 reactors under construction. Similarly, China is adding photovoltaic solar capacity at a rate of over 200 gW per year, giving it an annual rate of increase greater than the total installed solar capacity of either the US or the EU (Bradsher, 2024).

Most people who follow these issues have an opinion about the justifiability of these policy decisions. What is often lacking, however, is a principled basis for those judgments. Ideally, we would like a fully articulated theory of justice that could be applied to generate an assessment of the tradeoffs involved in any climate change policy, in order to determine the correct course of action. This would give us a principled basis for deciding whether China did the right thing back in 1996—and by extension, whether India and many other developing countries are currently doing the right thing in pursuing what amounts to the same policy of prioritizing growth over carbon abatement. It has proven unfortunately difficult to develop such a theory.[5] Much of this is due to the temporal structure of the climate change problem, arising from the fact that most of the damages that will result from present emissions occur in the relatively distant future. The application of simple tort principles, for example, is complicated by the fact that, by the time significant harms occur, most of those who are responsible for producing the relevant emissions will be dead.[6] The application of principles of distributive justice is complicated by the fact that one cannot redistribute goods back and forth between people in non-contemporaneous generations. Many normative theories that provide sensible guidance when dealing with contemporaneous moral problems generate absurd or deeply counterintuitive implications when applied to questions involving the distant future, or relations between existing people and those who have not yet been born.

Philosophers have been aware of this difficulty for a long time. John Rawls, writing in *A Theory of Justice*, observed that "The question of justice between generations… subjects any ethical theory to *severe if not impossible* tests" (before going on to present a view that he subsequently retracted and revised several times) (Rawls, 1971, 284). Derek Parfit identified not one but two profoundly paradoxical outcomes that seem to follow from elementary consequentialist principles when applied to moral problems involving future persons.[7] These difficulties were considered mainly a curiosity at the time, but climate change has conferred much greater urgency on the issue. Thus, philosophers have found themselves struggling to find something coherent, never mind sensible, to say about the various normative questions that arise inevitably in the formulation of climate change policy.[8] My objective in this paper will be to provide an overview of these challenges, since it is only

[3]Data is from World Bank, expressed in 2022 U.S. dollars: https://data.worldbank.org/indicator/NY.GDP.PCAP.CD?locations=CN (accessed March 2, 2026). Real per capita means adjusted for inflation, which implies that the average Chinese citizen in 2022 was 17 times richer than in 1996.

[4]Over the entire period, declining death rates resulted in approximately 8 million babies surviving who would otherwise have died. Data from World Bank: https://data.worldbank.org/indicator/SP.DYN.IMRT.IN?end=2020&locations=CN&start=1996 (accessed March 2, 2026).

[5]As Idil Boran writes, "If one is hard put to find a unified conception of climate justice, it's because there isn't one" (Boran, 2019, 27).

[6]Many theorists attempt to avoid the difficulty by assigning responsibility to group agents, such as nations (von Allmen, 2022). This generates its own difficulties (Posner & Weisbach, 2010, 101–105).

[7]I am referring to the "non-identity problem" (Parfit, 2017) and the "repugnant conclusion" (Parfit, 1982).

[8]Both Jamieson (2014) and Gardiner (2011) illustrate the difficulty.

by seeing the difficulty of the problem that one can understand why many of the positions that end up being taken by the more sophisticated participants in these debates seem initially so implausible.

My strategy for organizing the discussion will be to divide normative ethical theories along the conventional lines of consequentialist and deontological approaches. I will then show how both schools of thought have struggled to provide a coherent framework for analyzing the issues involved in climate change and intergenerational justice. One of the standard objections to consequentialism, especially in its utilitarian variants, is that it suffers from a "demandingness" problem (which is to say, it generates an excessively onerous set of moral obligations). The extension of consequentialist principles to include future people risks supercharging this demandingness problem, generating the problem of what I shall refer to as *overdemandingness*. On the deontological side, the problem has the exact opposite structure. By eschewing maximization, many deontologists find themselves treating most of the benefits that we provide to future generations as supererogatory. This generates an *underdemandingness* problem when it comes to the question of climate change, because it becomes far too easy to compensate future generations for the costs we are imposing on them, which then undermines the obligation to engage in any climate change mitigation.

In order to see how these twin problems wind up structuring the discussion, I will show how a simplistic application of either approach generates a straightforwardly absurd conclusion (which everyone recognizes and seeks to avoid), a somewhat absurd conclusion (which only extremists embrace) and a defensible but very controversial conclusion (which many wind up accepting). I will end by showing how this dynamic pushes climate ethicists toward what I will call the "narrow middle ground," of views that are more plausible. This results in two robust positions, namely prioritarianism for consequentialists and Coasian bargaining for deontologists. Neither view is without problems, but both have the merit of avoiding the extremes of overdemandingness and underdemandingness.

1. Consequentialist Approaches

It is difficult to make the case that the introduction of carbon dioxide into the atmosphere is intrinsically wrong, or that there could be a deontic prohibition against it, since every time we exhale we are making a contribution to the stock of atmospheric CO_2. Carbon emissions are only a problem because humans are changing the equilibrium of the atmospheric carbon cycle by digging up deeply buried hydrocarbons and burning them. This previously sequestered carbon persists in the atmosphere for millennia in the form of CO_2, which allows sunlight to pass but absorbs infrared radiation emanating from the earth, thereby contributing to a global warming process that is, by geological standards, extremely rapid. This CO_2 production occurs as a byproduct of activities that are largely beneficial, and so thinking about the problem of climate change almost necessarily requires balancing costs against benefits. Because of this, many theorists have found some form of consequentialism to be a natural framework for evaluating our options and determining our obligations (Broome, 1992, 2013).

A second major attraction of the consequentialist framework is that it appears to offer a simple way to handle the temporal dimension of the climate change problem, by drawing on an analogy between time and space. One of the central objectives of Peter Singer's influential article, "Famine, Affluence and Morality" (Singer, 1972) was to show that spatial distance between individuals is irrelevant from the moral point of view. If one is in a position to confer a significant benefit on some other person, at lesser or minimal cost to oneself, then it does not matter where that person is located in space—directly in front of one's eyes, or on the other side of the planet—one has the same obligation to confer the benefit. Many have thought that this argument generalizes easily from space to time. Parfit, for example, presented a now-widely-cited argument aimed at showing that the introduction of *delay* between an action and its effects should not change the moral calculus that

determines the permissibility of the action (Parfit, 1984, 356).[9] This has tempted many to draw a parallel between "location in space" and "location in time," and to think that both are irrelevant to the goodness or badness of events, and by implication, to the permissibility of the actions that bring them about.

Another feature of Singer's article, however, was that it brought to the fore a feature of consequentialist reasoning that many have found troublesome, which is that it risks allowing moral considerations to override self-interest too completely.[10] If an action is morally permissible only when it brings about the best outcome overall, and there are millions of people suffering various forms of abject deprivation, whose suffering has the exact same moral urgency as if it were happening before our own eyes, then we will seldom be entitled to consume more than a small fraction of our income, much less indulge in such unnecessary pleasures as buying a cup of coffee, watching a movie or playing a video game. This is the demandingness problem. It has generated a large philosophical literature, which there is no need to survey here (Kagan, 1984). The important point, for our purposes, is that the expansion of the circle of moral concern to include, not just all living persons (and possibly animals), no matter how far removed from us in space, but all future persons (and possibly animals) as well, no matter how far removed from us in time, has the potential to radically exacerbate the demandingness problem. (For simplicity I will avoid dealing with the special problems raised by the inclusion of dead people—or people "living in the past"—in the circle of moral concern, because time's arrow makes it difficult either to benefit or harm them.)

The problem is not just that future people (however we might choose to identify them) vastly outnumber us. The more serious problem stems from the fact that they are, in many respects, much easier to help than existing people, largely due to the productivity of *investment*.[11] One could give a man a fish and feed him for a day, teach him to fish and feed him for a lifetime or else construct and stock a trout pond and feed his whole family in perpetuity. Resources that are used to produce consumption goods can almost always be redirected to the production of capital goods in a way that would generate a greater sum of satisfaction over time. This means that as long as there is a positive rate of return on capital investment, saving will be preferable to consumption. Depending on the exact framing, this can have a number of more or less perplexing implications:

1.a. The straightforwardly absurd conclusion

The most obvious objection to treating future events and present events as having equal moral importance is that our concern over future events seems closely related to the fact that they will someday *become* present ones. There is enormous potential for confusion lurking about in expressions such as "your action imposes a harm on a future person." Although the word "future" in this sentence looks like an ordinary adjective picking out a particular type of person, it is better understood as indicating a verb tense. In other words, the way to interpret the expression is along the lines of "your action *will impose* a harm on some person, in the future." The wrongness of the action arises from the future effect on an actual person, not the present effect on a future person. This suggests that we do not care about future people *per se*, we care about them only insofar as they can be expected to turn into actual people. (It also explains why we do not care about possible people *per se*, we only care about the ones that can be expected to turn into actual people.)

There is an interesting analogy to be drawn here to debates between presentism and eternalism in ontology. Upon reflection, most people seem committed to a type of moral presentism, which implies that future harms are not really that bad so long as they remain in the future (that is, so long as they have not happened yet). Because of this moral presentism, expanding the circle of moral concern to include future people gives rise to a number of paradoxes involving our attitude toward

[9]For the influence of this argument, see Gosseries and Meyer (2009).

[10]A concern expressed influentially by Bernard Williams (Smart & Williams 2013).

[11]Future people are, in this respect, a type of "utility monster" (Nozick, 1971, 41). See Stanczyk (2021, 303).

future goods and bads that remain confined to the future without ever becoming actual. Suppose, for example, there is a bad event that is scheduled to occur in the future, but one has the ability to postpone the day of reckoning indefinitely. Suppose further that each time one opts to delay it, the choice increases somewhat the projected badness of the event. Simple "expanding the circle" arguments suggest that one is wrong to delay it, because that would involve elevating the interests of present people above those of future people.

Many of these paradoxes have been discovered and rediscovered so many times that it is difficult to trace their origin (e.g. Arntzenius et al., 2004; Pollock, 1983). An early version can be found in Lewis Carroll's *Through the Looking Glass*, where the Red Queen offers Alice jam "every other day," which unfortunately no one has ever had a chance to eat, because "the rule is, jam to-morrow and jam yesterday—but never jam to-day" (Carroll, 1981, 163). The "jam tomorrow" expression was picked up by John Maynard Keynes, in "Economic Possibilities for our Grandchildren" (Keynes, 2010, 370), in order to discourage excessive high-mindedness with respect to saving. From this point, the paradox became well-known among economists, and was given its canonical articulation in 1960 by Tjalling Koopmans, who formulated what he called the "paradox of the indefinitely postponed splurge" (Koopmans, 1960). Although accompanied by some rather imposing mathematical apparatus, the basic paradox can be stated simply: What if there were an investment opportunity, such that the choice was always between cashing it in for some quantum of satisfaction in the present or waiting a bit longer in order to receive a slightly larger quantum of satisfaction in the future? Absent some form of preference for present over future satisfaction, a person or society might wind up always choosing the larger, later satisfaction, and therefore wind up never achieving any actual satisfaction. This would be a case of always choosing jam tomorrow over jam today.

Following this line of reasoning, Henry Sidgwick's counsel "that a smaller present good is not to be preferred to a greater future good" (Sidgwick, 1981, 381), while initially sounding like a sensible principle of prudence, turns out to generate straightforwardly absurd conclusions when applied without qualification.[12] For a principle to require total sacrifice of all real satisfaction, for the entire foreseeable future, seems *overdemanding*. Sophisticated consequentialists have therefore sought to avoid formulating their view in a way that will generate this implication (Rendall, 2018). As we shall see, however, they have largely resisted Koopmans' recommended solution, which was to introduce a *time preference* into the social welfare function.

1.b. The somewhat absurd conclusion

Another problem with treating space and time as strictly analogous is that the number of individuals distributed through space is finite, whereas the number of future people is potentially infinite. As a result, expanding the circle of moral concern to include future people has the potential to introduce all of the well-known problems associated with the assessment of infinite utility streams into the consequentialist calculus. Apart from the fact that such streams cannot be compared to one another through simple aggregation, and that they undermine probabilistic calculations, infinite utility streams also have the unfortunate quality of having larger aggregate value than any finite sum. As a result, it is easy to construct scenarios in which it would appear justified to impose enormous costs in the present, in order to generate a stream of trivial benefits extending out infinitely into the future. For example, when people talk about species preservation, the stated objective is usually not to postpone the date of eventual extinction, but rather to preserve the species forever. Unless this is subject to some further qualification, this implies that we should be willing to accept arbitrarily large sacrifices in the present (e.g. most of world GDP) in order to protect a single species.

[12]Note that he does allow for adjustments based on risk, which allows one to handle certain cases, but does not avoid the basic problem.

This difficulty has figured centrally in the debates over *just savings* among economists. Kenneth Arrow gave the following, somewhat droll formulation of the basic problem:

> Imagine that an investment opportunity occurs, available only to the first generation. For each unit sacrificed by them, a perpetual stream of α per unit time is generated… Thus, any sacrifice by the first generation is good. Strictly speaking, we cannot say that the first generation should sacrifice everything, if marginal utility approaches infinity as consumption approaches zero. But, we can say that given any investment, short of the entire income, a still greater investment would be preferred (Arrow, 1999, 14).

He concludes that treating all costs and benefits alike, regardless of when they occur in time, is overdemanding ("the strong ethical requirement that all generations be treated alike, itself reasonable, contradicts a very strong intuition that it is not morally acceptable to demand excessively high savings rates of any one generation, or even of every generation" [Arrow, 1999, 16]). Beyond the issue of savings and investment, the problem also shows up in particularly acute form when dealing with questions involving research funding, because important discoveries, once made, tend to become available "for free" to all subsequent generations, and so produce an open-ended and often massive stream of benefits. (For example, the value of the mRNA vaccine development platform, which was deployed successfully in combating the COVID-19 pandemic, confers a benefit on future generations that is, if not infinite, then at least immeasurably large.) This can generate an overdemandingness problem when it comes to considering how much sacrifice present generations should be willing to accept in order to fund innovation, because the production of infinite utility streams in this domain often seems quite realistic. It is not difficult to see how climate change calculations become vulnerable to similar difficulties. Every megaton of CO_2 that is not released into the atmosphere creates an open-ended stream of future benefits. One would not want to conclude, however, that we should be willing to make arbitrarily large sacrifices in the present in order to achieve trivial improvements in the trajectory of global warming. Simplistic "expanding the circle" reasoning to include future people, however, suggests that we may not be entitled to produce the emissions required to feed, clothe and house existing human populations.

1.c. The defensible but very controversial conclusion

Every aggregative social welfare function must have some way of rendering commensurable costs and benefits that occur at different points in time. The standard approach among economists is to introduce a discount function, which reduces the weight given to an outcome the further removed it is from the present. Both Koopmans and Arrow sought to resolve the overdemandingness problem by introducing a *pure time preference* into the discount function, so that there would always be at least a slight social preference for present over future satisfaction.[13] In other words, they both argued that society should assign future satisfaction a lower present value the further removed it is in time, *merely* because it is further removed. (It is important to be careful in formulating this thesis: time preference is often misdescribed as the view that present satisfaction should be assigned greater value than future satisfaction. This is not quite correct. It is not the actual value of future satisfaction that is reduced, only its *present* value, so long as it remains in the future.[14])

[13]It has sometimes been suggested that the introduction of a discount rate does not really solve the paradox of the indefinitely postponed splurge because if the future benefits are growing as a sufficiently high rate, they may overwhelm the discount rate, and so continue to generate the paradox (e.g. Tarsney, 2017). The difference is that with a discount rate of zero, the benefits need only be growing incrementally, whereas with any positive rate, the benefits must be growing exponentially, in order to generate the paradox.

[14]As Kian Mintz-Woo (2019), 103) has observed, this confusion between present and future value is similar to the confusion experienced by some between expected value and outcome value (or payoff) in decision theory.

Many economists, one is inclined to suspect, have accepted this position simply because it gets rid of all the unbounded functions in their models and so makes the optimization problem mathematically tractable. Philosophical consequentialists, by contrast, have been extremely resistant to the suggestion that a time preference could be a legitimate component of moral decision-making (primarily for Singer-Parfit reasons) (Gosseries & Meyer, 2009). And yet the attractions are obvious, given the need to avoid the various problems arising from a failure to discount future satisfaction. Naturally one does not want to set the rate too high, otherwise it may generate an equally problematic outcome, in which action to combat climate change can be postponed forever (on the grounds that future people will always be better equipped to tackle the problem than we are). Thus the compromise position that has often emerged involves adopting a very low social discount rate, based on factors other than time preference.

The most obvious candidate for this is risk. In articulating his principle of prudence, Sidgwick always included the qualification that a smaller-sooner satisfaction might reasonably be preferred over a larger-later one if enjoyment of the former was more certain than the latter.[15] The major basis for this in the case of individual decision-making is the probability of death, which makes it rational to take some fraction of one's pleasures now rather than always postponing. While assessing decisions from the "point of view of the universe," as Sidgwick recommended, precludes such concerns, many consequentialists have argued that the probability of human extinction should function in the social calculus in the same way the probability of death does for individuals (Stern 2014).[16]

Modulo certain views about nuclear proliferation and hostile artificial intelligence, this move introduces a discount rate that remains uncomfortably close to zero. This creates a temptation to appeal to *economic growth* as grounds for increasing the discount rate somewhat further (Broome, 1994). Because of the productivity of investment, consuming something now genuinely is more costly than consuming something in the future. Because $100 invested at 6% interest yields $200 in 12 years, a $200 loss 12 years from now should only count for $100 in the present. This intuition received formal expression in the so-called Ramsey formula for discounting, which explicitly incorporates growth, risk and time preference. Consequentialists have been inclined to adopt what one might think of as the 2/3rds Ramsey formula, which is to keep the growth and risk, while leaving out time preference (Stern 2014).

One can see the force of these considerations even among so-called longtermists (whose motivating philosophical framework is usually consequentialist and so generates the usual dilemmas around discounting). If the measure of how much one cares about something is the amount that one is willing to sacrifice for it, then caring about the welfare of people in the very distant future entails a greater willingness to sacrifice the welfare of those in the present. As a result, longtermism takes the most unattractive feature of utilitarianism—namely the willingness to trade off satisfactions from one person to another willy-nilly—and supercharges it, generating a willingness to trade off the interests of present people against even those in the very, very distant future (Cowen, 1992). This leaves proponents of the view vulnerable to charges of misanthropy (especially if their moral framework is one that imposes no special constraints on trading off lives against one another). Like the well-known figure of the socialist who loves "the people," just not any particular person, many longtermists seem to love "humanity," just not the humans who happen to be alive at the moment.

Perhaps in order to avoid this difficulty, William MacAskill makes the rather surprising admission in his widely read introduction to longtermist thought: "I'm not claiming that the interests of present and future people should always and everywhere be given equal weight. I'm just

[15]After the statement of his principle of prudence, Sidgwick added the qualification "(allowing for differences of certainty)" (Sidgwick, 1981, 381).

[16]For an account of why risk should be incorporated into the discount rate specifically, see Mintz-Woo (2019).

claiming that future people matter significantly. Just as caring more about our children doesn't mean ignoring the interests of strangers, caring more about our contemporaries doesn't mean ignoring the interests of our descendants" (MacAskill, 2022, 11). This is a rather significant concession. Indeed, it raises questions about whether one can hold this view and still claim to be a longtermist. After all, no one is arguing that we are entitled to ignore *completely* the interests of our descendants. And yet even the introduction of a very small discount rate, such as 0.1%, precludes any practical concern over events occurring beyond the present millennium (although it never reduces the level of concern to zero).

And yet even though adopting a small positive discount rate allows consequentialists to avoid absurdity, it generates a moral framework that remains extremely demanding, possibly overdemanding. With respect to climate change, for example, it implies that we should be prepared to make surprisingly large sacrifices in order to help people in the future who are expected to be comparatively quite well-off—a great deal better off than we are. Indeed, this was one of the central objections to the recommendations of the Stern Review on the Economics of Climate Change, which was commissioned by the UK government in 2006 (Nordhaus, 2007). The chief architect of the review, Nicholas Stern, was persuaded by the moral case against time preference and so employed a discount rate based solely on the risk of human extinction and the expected long-run growth of the global economy, which led him to adopt a relatively low rate of 1.3%. This was the major factor driving his conclusion that the UK should adopt a carbon price approximately 10 times higher than the mainstream recommendations being put forward at the time.

Critics of the report argued that these recommendations flowed from a view on intertemporal tradeoffs that seemed overdemanding with respect to the present. Christian Gollier, for example, observed that Stern's discount rate implied "that one should be ready to give up as much as 60% of GDP per capita, now and forever, to raise the growth rate of GDP over the next two centuries from 1% to 2% per year" (Gollier, 2006, 4). Furthermore, if one were to apply the same discount rate to the evaluation of policy options in other domains, the results would be quite radical. At the very least, it suggests that we should be saving much, much more than we are—perhaps upwards of 90% of GDP. Indeed, if one considers the overall value of the bequest that we are leaving to future generations, we do not seem to be making even a half-hearted attempt to maximize its value. The current division of spending between health care and health research, for example, suggests that we have our priorities almost exactly backward—rather than spending 95% or more on care, we should probably be spending that much on research.

It is important to emphasize that Stern's view, as an instance of the more general strategy of introducing a small discount rate that excludes time preference, is *defensible*. But it is not in reflective equilibrium with many of our practices, or even current understandings of what economic justice requires. This raises the question whether we are obliged to embrace such an exotic and revisionist conception of morality in order to articulate what is wrong with unrestricted greenhouse gas emissions. Specifically, do we really need to adopt a moral theory that forces us to change our considered view on questions such as just savings, health care spending and myriad other policy issues, in order to formulate a reasonable response to climate change? Many will find it difficult to believe that correcting what is, after all, a mere pollution problem could have such wide-ranging implications for other our moral commitments. So while the low-discount rate position is defensible, it seems also to be unnecessarily controversial.

2. Deontological Approaches

The larger problem for consequentialist approaches to thinking about climate change is that many people are simply unpersuaded by the central premise, that we are morally committed to *maximizing* benefits to future generations (Stanczyk, 2021, 296). Consequentialists appear to be arriving at a reasonable conclusion—that we should do something about climate change—for the wrong reasons. It is of course true that climate change represents a clear-cut failure of maximization.

There is very broad scientific agreement that, if nothing is done to curtail greenhouse gas emissions, future generations will be substantially worse off than they *might have been*, in a world without climate change (Bilal & Känzig, 2026; Burke et al., 2015; Waidelich et al., 2024). But they are still likely to be much better off than we are now, in multiple dimensions. The fact that they could have been better off still does not seem to involve any serious dereliction of duty on our part. For instance, very few people seem to fault past generations for their rather conspicuous failure to maximize *our* well-being. Commentators such as Olúfẹ́mi Táíwò (2022), whose work on climate justice bristles with indignation toward past generations over slavery and colonialism, appears totally untroubled by their low savings rates. This indifference to past violations of the supposed obligation to maximize future well-being suggests that we do not really take ourselves to be under this obligation either.

Whether it is ultimately defensible or not, many philosophers have the intuition that our obligations toward future generations amount to leaving them no worse off than citizens of moderately affluent countries are today (Barry, 1999). Rawls, for example, when considering the question of just savings, suggests that satisfying our obligations toward future generations requires passing along no more than a steady-state economy, which replenishes its capital stock but does not engage in investment aimed at expanding production. He argues therefore that present generations have an obligation to save only until per capita GDP is high enough to maintain just institutions (i.e. so that society is no longer "burdened"), after which point all savings become discretionary. (It is perhaps worth noting that in 1971, when Rawls published this view, he took the United States to be well beyond that point. Per capita GDP in the United States has more than doubled since that time, which implies that it could be drawn down quite considerably without actually burdening that society.[17]) This analysis is typically motivated by a commitment to equality, or to a somewhat more diffuse conception of fairness, which generates the requirement that all generations find themselves similarly situated and that the consumption of any one generation not occur at the expense of later generations.

Some have been tempted to interpret this intuition as a constraint on taking actions that will harm future generations (Davidson, 2008; Shue, 1999). The problem with this view, however, when applied to climate change, is that a large fraction of greenhouse gas emissions are generated as a byproduct of activities that also benefit future generations, and so a simple prohibition on harm winds up being either gridlocking (if construed strictly) or underdemanding (if construed more flexibly). To see the potential for gridlock, consider the example of China's reduction in infant mortality during the period since the Kyoto Accord. Suppose that China had adopted carbon abatement policies in 1996 that resulted in a somewhat lower rate of growth (e.g. opening fewer coal-powered electricity generation plants). Suppose also that this resulted in a somewhat lower rate of decline in infant mortality. Back-of-the-envelope calculations suggest that even a 10% reduction in the rate of decline would have resulted in almost 800,000 more babies dying over the period 1996–2020. As a result, a policy of building fewer coal-fired power plants would clearly have harmed future generations, just as the actual policy, of building more coal-fired power plants, is destined to harm future generations. A strict prohibition on harm (e.g. acting in a way that will foreseeably result in the death of a person not yet born) therefore excludes both action and inaction.

Because of the inevitability of tradeoffs such as this, deontological approaches to climate change typically focus on the *net* effect of our actions on future generations. This approach, however, risks being underdemanding because the expected losses due to climate change are, in most probable near-term scenarios, dwarfed by the benefits produced through ordinary processes of economic growth, not to mention technological change. Naturally, if climate change were to run unchecked for centuries, there would eventually come a point at which the damages became so severe that they would depress living standards below those of the present day. Similarly, there are some outlying catastrophe scenarios in which climate change sets off a feedback process that results in sudden,

[17]Date from World Bank: https://data.worldbank.org/indicator/NY.GDP.PCAP.KN?locations=US (accessed March 2, 2026).

uncontrolled warming. The normative challenge arises in the medium term, with the most probable scenarios, because economic growth is such a rapid force for improvement, particularly in developing countries, that it easily outweighs the damages caused by the more slow-moving process of climate change. Thus a simple prohibition on harm is far too permissive because future generations wind up being net beneficiaries of our actions under all but the most extreme scenarios.

Furthermore, because growth is largely a byproduct of self-interested action, not moral obligation, future generations can be expected to benefit from economic growth even if they are not entitled to it and we are not obliged to produce it. This means that, in principle, the fruits of this growth can serve as an enormous compensation fund to make up for costs that we are imposing upon them. For example, at the current rate of sustained economic growth, the size of the global economy is doubling every 25 years. Suppose that, as a consequence of unchecked climate change, people 25 years from now are obliged to dedicate as much as one half of their economy to a combination of carbon scrubbing and climate change adaptation (e.g. sea wall construction, extreme weather protection, etc.). From a steady-state perspective, this need not be unjust, since the resources they will be dedicating to these pursuits are not ones that they were entitled to in the first place. This seems underdemanding, insofar as it entitles us to do nothing about climate change.

Some scholars, it should be noted, have fallen prey to misinformation on this score, and so believe that climate change literally poses an "existential" risk to the human species, or that within the next few decades it is likely to depress living standards below what they are today (Hamilton 2011; Mulgan, 2011; Wallace-Wells 2018). The latter belief is often based on a misunderstanding of the scientific literature, which too often speaks about "loss of GDP" in 2050, or 2100, without making clear that this loss is relative to what GDP *might have been* at those future dates, not relative to what it is today.[18] It would, of course, generate an enormous simplification of the normative problem if our actions were poised to transform the planet into an uninhabitable hellscape. In such a case, the cost of burning fossil fuels would easily outweigh the benefits, and so a straightforward prohibition on burning them would be in order. The problem we face, however, is that we know the current level of fossil fuel consumption is too high, but we also know that reducing consumption to zero would be much too low. The task of a normative theory is to determine what level of consumption is justifiable, given the important benefits that fossil fuels provide but also the costs that they entail. This sort of weighing of costs and benefits is obviously not the strong suit of deontological theories of justice. But in addition to this, egalitarian theories encounter special problems when extended to intergenerational relations.

2.a. The straightforwardly absurd conclusion

Philosophers who are inclined to think about intergenerational relations in terms of principles, rather than consequences, realized very early on that the combination of egalitarianism and "expanding the circle" to include future people generated conclusions that were absurd. For example, critics of Rawls's characterization of distributive justice noted that a straightforward extension of the original position thought-experiment generated what we might call the "eternal poverty *reductio*." At first glance, it seems as though the correct way for Rawls to have handled questions of intergenerational justice would have been to include in the original position past, present and future individuals. He does, after all, characterize the basic structure as instituting "a fair system of cooperation between generations over time" (Rawls, 2001, 160). It stands to reason that every generation should be represented in the original position, with fairness then guaranteed by a veil of ignorance that deprives individuals of knowledge (*inter alia*) of which generation they belong to.

[18]For discussion, see O'Neill (2023) and Ritchie (2024).

The problem with this approach to intergenerational relations is that it leads to the conclusion that savings are not only non-obligatory but also actually impermissible.[19] If there are any productive investment opportunities available, savings will generate growth, which will make future generations better off than present ones. This makes the present generation the worst off, which results in the difference principle assigning *lexical* priority to the primary goods endowment of this generation. This requires, in turn, maximizing present consumption, which precludes savings. As Robert Solow observed:

> This result follows from the basic principle itself. Capital could be accumulated and consumption increased subsequently, but only at the cost of a lower standard of living for earlier generations. It is part of Rawls's general argument for the max-min criterion that we should regard earlier and later generations as facing each other contemporaneously when the social contract is being drawn up. But then it is hardly surprising that the preferred strategy refuses to make some people poorer than others in order to make the others richer, just because the first group can be given the essentially arbitrary label of "earlier" (Solow, 1974, 36).

Thankfully Rawls realized this was a problem before *A Theory of Justice* went to print and so refrained from simply expanding the original position to include individuals situated at different points in time. The alternative that he presented, however, struck many as unmotivated, and Rawls himself revised his view in later work.[20] It is not difficult to see why. Any conception of justice that assigns significant weight to the principle of equality is likely to regard economic growth as a source of unjust advantage for future generations. Furthermore, unlike natural inequality, intergenerational inequality is clearly the product of human social institutions, because it is a consequence of economic growth. If the normative framework tolerates leveling down, this can generate the implication that we should not be doing anything to improve the circumstances of later generations. There should be no progress, on this view, because it would be unfair to those who had the bad luck to be born in earlier periods of human history. Obviously, this is absurd. Apart from condemning any improvement in the human condition, it suggests that, to the extent that we are unable to avoid helping our descendants in some way, we should also impose some harms on them, in order to ensure that equality remains satisfied. This has the unfortunate implication that if environmental damage mitigates some of the unjust advantages enjoyed by later generations—produced not just through economic growth, but also through the advancement of knowledge—it may be required.

2.b. The somewhat absurd conclusion

Another way of approaching the problem is to suppose that the general egalitarian constraint takes a less restrictive form when applied to intergenerational relations. The standard way of weakening egalitarianism is to transform it into sufficientarianism, or the view that there is a threshold below which individuals should not be allowed to fall, but that any differences above and beyond this threshold are not an affront to justice. There is a very large family of views that interpret justice as imposing the requirement that all persons enjoy a specified bundle of "basic goods," or an expansive conception of "human rights" or even just a minimum income level. According to such views, there is nothing wrong with future generations being richer than us, and so the eternal poverty *reductio* is avoided. The requirement is simply that we not engage in actions that will foreseeably push future generations below the threshold of sufficiency (or violate their social rights, deprive them of basic goods, etc.)

[19]Some have considered this not self-evidently wrong (e.g. Gaspart & Gosseries, 2007).

[20]Rawls revised his view significantly between the first and second edition of *A Theory of Justice*. He also articulated a different view in Rawls (2001). For discussion see Karnein (2022).

Closely related to these views are proposals that appeal to the concept of *sustainability* as the most important normative principle governing intergenerational relations. Although it has been given various formulations, sustainability is typically articulated in a way that makes it essentially sufficientarian. There is, for example, the well-known Brundtland definition, which defines sustainable development as "development that meets the needs of the present without compromising the ability of future generations to meet their own needs" (Brundtland, 1987). One can find other definitions that require present generations to refrain from doing anything that will deny future generations the opportunity to achieve the same level of utility, or satisfaction of their wants, as we currently enjoy (Barry, 1999).

A certain amount of complexity is introduced into these definitions by the recognition that values may change over time, and so the needs and desires of future generations may be quite different from our own. Thus the formulation of these principles often uses quite open language, in order to avoid imposing on future generations our own, potentially parochial views about what is more or less important. This serves to focus attention on the net value of the inheritance, with the threshold of sufficiency specified in some relatively abstract currency. The consequence is that within the bequest that we make to future generations, these views permit practically unlimited substitution of goods from one category to another. This is obvious in the case of models that focus on GDP, or that use money as a metric for calculating damages. But it is equally true of more explicitly deontological approaches. Simon Caney, for example, uses the language of human rights to articulate what amounts to a sufficientarian view, according to which we are prohibited from depriving future persons of several goods, including their means of subsistence (Caney, 2010). This requires us to refrain from doing anything that will compromise the output of agricultural systems in the future. At the same time, the constraint is consistent with large-scale loss of farm land (e.g. to flooding), because these losses can be made up for by improvements in other areas, such as the development of higher-yield crops or improved synthetic fertilizers. As a result, no one need be deprived of the means of subsistence, so long as losses in one area can be compensated for by gains in another.

Thus the problem with the sustainability framework, as many environmentalists have complained, is that it is far too permissive, especially when it comes to problems like climate change. We produce so many benefits for future generations, as an incidental byproduct of activities aimed at improving our own well-being, that the *net* value of their inheritance is projected to be positive under all but the most extreme scenarios. This is obvious with respect to GDP, but is equally true with respect to many other dimensions of well-being.[21] Consider the malaria vaccine that is currently in the early stages of deployment, which appears to be about 80% effective. Currently, there are over 240 million malaria cases per year, resulting in approximately 600,000 deaths, and so just this one vaccine stands poised to save a half-million lives per year. If one factors in all of the other contributions that we are making to human health and life expectancy—many of which are simply a byproduct of economic growth—these easily make up for the damages caused by climate change, at least between now and the end of the century. Thus a principle like Caney's "human right to life," applied to future generations, is far too easily satisfied.[22] Of course, the benefits to future generations could be even greater if we were to do something to mitigate climate change, but if we are not concerned with maximization, only with ensuring that no one falls below a standard of sufficiency, then it is difficult to motivate much concern over these damages. Thus the combination

[21]Note that GDP impacts of climate change are often presented in a misleading way, which suggest that GDP is likely to be depressed below present levels. For discussion see Heath (2021, 83–84).

[22]I am assuming that the right is to be construed probabilistically. Climate change increases the chances of dying from flooding or extreme weather. Other actions we take may reduce the chances of dying from malaria, malnutrition, earthquakes, communicable diseases, and so forth Presumably, the right to life is satisfied if, on balance, net chances of dying are not increased. If the right is interpreted strictly—that no one is permitted to do anything that increases anyone's chances of dying, then it cannot be satisfied.

of sufficientarianism and substitutability generates the somewhat absurd conclusion that we are not obliged to do anything about climate change, or at least not between now and the end of the century.

2.c. The defensible but very controversial conclusion

If the underdemandingness problem follows from the joint commitment to sufficientarianism and substitutability, the obvious way to salvage sufficientarianism is to abandon substitutability. Thus it has become conventional among environmental philosophers to distinguish between weak sustainability, which permits substitution of social for natural capital within the bequest made to future generations, and strong sustainability, which does not (Neumayer, 2010; Norton, 2005). The intuition, stated coarsely, is that if we destroy various aspects of the natural environment, we cannot make it up to future generations by giving them more consumer goods. This is taken to address the underdemandingness problem, because it imposes an obligation to pass on to future generations a natural environment that meets certain sufficiency conditions, regardless of what other benefits we may be producing.

Unfortunately, while the basic idea can easily be stated, it turns out to be a great deal more difficult to articulate in precise terms. In standard definitions of sustainability, the feature that opens the door to substitution of social for natural capital is the specification of the sufficiency threshold in terms of human impacts (e.g. in the Brundtland definition, the concern is that we not compromise the ability of future generations to meet their *needs*). This is combined with the rather natural supposition that the adequacy of the bequest made to future generations must be assessed in terms of its net value. It is, after all, inevitable that the impacts we have on them will not be uniformly positive, but will involve some combination of goods and bads. But if one puts these two ideas together—concern with the net value of the bequest and specification of that value in terms of human impacts—then it follows that there can be no *normative* prohibition on the substitution of social for natural goods. It may turn out to be empirically difficult (e.g. extremely costly) to achieve such substitutions, but the weak sustainability theorist's claim is merely that any bequest that puts future generations above the sufficiency threshold is as good as any other, regardless of its specific contents.

Strong sustainability theorists have not always been clear on what would be entailed by their rejection of this framework (DesRoches, 2019). In many cases, weak sustainability theorists specify the sufficiency threshold in terms of human impacts, not because of anthropocentrism, but rather from a commitment to liberal neutrality. The concern, in other words, is that values change over time, and so not only is it difficult for us to know what future generations will want, it would be improper for us to foreclose the range of options available to them by locking in some particular set of commitments (Bykvist, 2007; Zwarthoed, 2015). This implies that we should be leaving them enough to meet "their needs" *according to their own conception*, not ours. This motivates a desire to leave them resources in a relatively indeterminate form, so that they can make substitutions as they see fit, depending on how their needs and desires develop. Thus a refusal to grant the permissibility of substitutions within the bequest not only constrains us but it also constrains future generations.

The problem this creates can be seen by considering the range of different categories that structure our bequest to future generations. These include, at very least: cultural artifacts, technological expertise, scientific knowledge, institutions, durable consumer goods, capital goods, elements of the built environment, natural resources, biodiversity, wilderness, ecosystems and dependent natural resources (Heath, 2021, 91–93). It is not difficult to construct cases in which it would be intuitively objectionable, with respect to widely shared first-order values, to make substitutions from any of these categories to some other. But turning this into a fully articulated account that specifies what or how much we owe to future generations in each category would rely upon a set of very precisely specified perfectionist values. Unless we feel that these values brook no disagreement, we might hesitate to impose them on future generations. So while adopting a strong sustainability position allows one to say, for instance, that present generations must adopt

non-depletionary policies with respect to each of the categories of natural bequest (ecosystems, biodiversity, etc.), because compensation for losses with gains in other categories is not permissible, it does so at the expense of adopting an extremely controversial normative framework.

Finally, the value judgments that speak in favor of a strong sustainability framework are a great deal less compelling in the case of climate change than with respect to many other environmental problems. Few people have the intuition that there is a correct temperature for the planet. We know that the range of natural variation has been quite significant on a geological timescale, and even a relatively proximate event, such as the end of the last ice age, involved warming of approximately 5 °C. The concern is really with the *effects* of rapid warming. Here one can see that the most damaging impacts of climate change are on human welfare rather than natural systems per se. The effects on human agriculture and the built environment, in particular, are confined to domains in which substitution seem permissible. For example, future generations can be compensated for many of the effects of extreme weather through investments in infrastructure, more demanding construction standards and relocation of certain populations. All of this may be extremely costly, but since the compensation all occurs within-category, strong sustainability offers no solution to the underdemandingness of the sufficientarian standard. Similarly, the creation of millions of climate refugees is not, strictly speaking, an environmental problem but rather a challenge to the adaptability of human social institutions. There are exceptions to this, such as ocean acidification and biodiversity loss, where compensation seems less plausible.[23] But in other major areas it remains a serious problem for the strong sustainability view.

So while one can see how proponents of a deontological stance wind up being pushed to adopt some form of strong sustainability commitment, in order to avoid the underdemandingness problems that afflict egalitarian and sufficientarian views, the position that results remains very controversial, primarily because it seems to require the abandonment of liberal neutrality. This may explain also why the view is relatively undeveloped in the literature. Most importantly, there has been no attempt to provide a detailed specification of the compensation arrangements that are and are not permissible within the total bequest we are making to future generations. There may be various reasons for this, but one of them is likely to be that the view becomes more controversial the more clearly it is specified, and so keeping it at the level of abstract pronouncements is a way of soft-pedaling its more unattractive perfectionist features.

3. The Narrow Middle Ground

This analysis makes it easier to see why many climate ethicists are exceedingly vague when it comes to articulating their normative commitments. It also explains why some have abandoned the hope of producing a normative theory that can answer fine-grained questions such as the correct carbon pricing level, or the appropriate balance of mitigation and adaptation. Idil Boran, for example, in *Political Theory and Global Climate Action*, argues the political philosophers should renounce the commitment to "axiomatic moral justification," whether in the style of John Rawls or Peter Singer (2018, 4). She recommends instead a process-based approach, focused on modest idealization of the deliberative practices prevailing between nation states and among civil society groups. Darrell Moellendorf, in *Mobilizing Hope* (Moellendorf, 2022), similarly recommends abandoning a foundational approach to thinking about climate justice. He instead accepts the 1.5 °C end-of-century warming target (which he describes as "a decent approximation of what we owe future generations" based on converging lines of argument) and focuses on issues that arise with respect to the attainment of that goal (Moellendorf, 2022, 22).

[23]It is worth noting that the biodiversity issue is handled separately at the international level, with the UN Biodiversity Conference meeting apart from the more highly publicized UNFCC climate change conferences.

Despite these pleas for theoretical modesty, some philosophers have nevertheless persisted in the attempt to adjust one or the other of the standard normative theories to make it capable of addressing the climate change challenge. Among consequentialists, the most compelling way of doing this is to adopt some form of *prioritarianism* (Adler et al., 2017; Adler & Treich, 2015; Ferranna, 2021). This can be thought of as a form of utilitarianism that has been modified to be more-or-less strongly biased in the direction of greater equality.[24] The bias is achieved by assigning greater weight to the welfare of individuals, in the aggregation function, the lower their absolute welfare level. As we have seen, egalitarian commitments tend to favor present over future generations because of the way that economic growth, not to mention advances in culture, science, and technology, benefits later generations. By introducing a modest egalitarian concern into the utilitarian calculus, prioritarianism is therefore able to assign greater weight to the interests of present over future generations, but without grounding this in any sort of *time* preference. If straightforward aggregation is overdemanding, and egalitarianism is underdemanding, it stands to reason that one can achieve the appropriate level of demandingness with a theory that assigns some weight to both aggregative and egalitarian concerns.

Technically, this can be accomplished within a utilitarian framework through modification of the social discount rate. The Ramsey formula, as mentioned earlier, contains three principal terms, reflecting the contribution of economic growth, risk and time preference. The growth term normally contains a modifier, the elasticity of marginal utility of consumption (η), which is introduced in order to correct for the fact that marginal utility declines as wealth increases. We get less "bang for the buck" from further economic growth, the higher the standard of living, but we also experience losses as less severe. Thus the Ramsey formula multiplies the expected rate of economic growth by the elasticity of marginal utility of consumption, then adds this to the combined measure of risk and time preference to give the discount rate. For example, if the expected growth rate is 2%, the incorporation of this raw number into the discount rate would lead contemporary planners to treat a loss of a given magnitude next year as 2% less important than a present loss of the same magnitude. But if η takes the (typical) value of 1.5, this means that individuals in the future, due to their increased wealth, will be less responsive to both losses and gains than they are in the present, and so contemporary planners should treat the loss next year as 3% less important than a loss this year. (It is mildly counterintuitive that the higher these numbers the *less* we should care about the future.)

It has often been observed, however, that the value of η can just as easily be used to reflect a normative judgment, about how much *we* care about gains and losses to individuals at different levels of wealth (Buchholz & Schumacher 2010; Groom & Maddison Pr, 2019).[25] It can be used as a measure of inequality aversion, with a higher number reflecting an egalitarian commitment (Dasgupta, 2008). The incorporation of a social discount rate with a higher value for η can be used to create a system of intergenerational prioritarianism (according to which present generations are given license to assign greater weight to their own interests because they are poorer than future generations).

From a technical point of view, there can be little doubt that this modification to the consequentialist approach accomplishes what it sets out to do. In effect, it avoids the overdemandingness problem by introducing a variable that functions like a slider, allowing consequentialists to adjust the demandingness of the normative theory until it is at the desired level. The question is whether this is just a work-around to the overdemandingness problem, or whether it corresponds to some deeply felt intuition about intergenerational relations. The question, in other words, is whether we really care about intergenerational inequality, or whether we are merely gerrymandering the

[24]Note that this is not how all consequentialist prioritarians think about it (e.g. Parfit, 2002).

[25]As Paul Kelleher observes (Kelleher, 2025, 28), the way that this point is articulated often invites confusion.

discount rate to produce a climate policy that seems appropriate relative to some other, unexpressed moral criterion.

On the deontological side, the most promising approach involves refraining from treating climate change as a general problem of intergenerational justice and approaching it instead as just a negative externality problem (e.g. comparable to sulfur dioxide emissions). As Ronald Coase (1960) famously observed, these externality problems can be decomposed into two quite separate questions: first, the question whether the externality should be produced, and if so at what level; and second, the question of who should bear the burden of adjustment of any departure from the status quo. The first is an efficiency question, the second a distributive justice question. The so-called Coase theorem is derived from his observation that (under certain idealizing conditions) the answer to the second question does not affect the answer to the first because regardless of who gets assigned the burden, the parties can be expected to contract to the outcome that achieves efficient production of the externality. Coase therefore favored an arrangement that, whenever feasible, would facilitate such negotiation. At the same time, he recognized that when transaction costs were too high, these negotiations would not occur and so the state would have no choice but to regulate (Coase, 1960, 17–18). In doing so, he argued, it should focus on determining the outcome that the parties *would have contracted to*, then imposing it through law. Ideally, this can be done in a way that also achieves the appropriate distribution of the burdens of adjustment, but if not, the allocation can always be adjusted through compensation.

Applied to the case of climate change, the temporal structure of the problem obviously precludes actual bargaining among the parties because many of those who will be affected have not been born yet. This leaves states with the challenge of determining what the outcome of such negotiations would have been, had the parties been able to meet. (Again, this question can be answered without having to determine the just allocation of the burdens of adjustment.) The crux of these negotiations comes down to how much those who are negatively affected by the externality would be willing to pay to get the emitter to stop producing it versus how much the emitter would be willing to pay to be able to continue producing it. The regulatory objective is then to prohibit the production of the externality wherever the former sum exceeds the latter. (Whether anyone *actually* gets paid to stop, or is compensated for damages, is the distributive justice question.) The efficient outcome can be achieved through the imposition of a carbon price, with the level set equal to our best estimate of the social cost of carbon (SCC) (Kelleher, 2025). This approach avoids the whole set of problems around substitutability that afflicts more general deontological approaches by setting aside the large-scale question of intergenerational justice and focusing just on the externality, trying to determine its efficient level. The objection to carbon emissions, from a Coasian perspective, is that it involves present generations unilaterally imposing costs on later generations, under terms that would not be agreed to if the parties had to come to a negotiated agreement.

Although when presented in this way the view does have a certain pleasing simplicity, many theorists find it counterintuitive because of the way that it attempts to treat climate change as a straightforward pollution problem, unconnected to any grand bargain between generations. In one respect this straightforwardness is illusory because the intergenerational problematic does recur. Calculating the SCC also requires a discount rate, and so the adoption of a Coasian perspective winds up importing into the deontological framework the central problem that consequentialists have struggled with. Of course, one might choose to think about the discounting problem in a way that breaks completely with the Ramsey approach (Heath, 2021, 261–264). At the same time, the fact that people in the future are very likely to be richer than us obviously affects the magnitude of the burden that is imposed upon them by the fact we are, in effect, forcing them to undertake adaptation in response to climate change. Thus one can also make a contractarian case for thinking about the social discount rate in the way that is codified in the Ramsey formula. This generates a significant narrowing of the difference between the Coasian and the prioritarian approach to the question. At the limit, the Coasian approach becomes just a way of providing deontological foundations for the policy prescription favored by consequentialists (carbon taxation, with a price

set to the best estimate of the SCC). The central advantage of the Coasian approach lies in the way that it focuses strictly on the externality problem, setting aside the general question of what we owe to future generations, as well as eschewing any commitment to the maximization of future welfare. Excessive carbon emissions, from the Coasian perspective, constitute a type of free-riding, and can be condemned strictly on those grounds.

4. Conclusion

It is not difficult to see why climate change poses a *pragmatic* challenge. It is a general feature of pollution problems that polluters have no interest in stopping what they are doing. Even when the outcome is inefficient, they still have a free-rider incentive to continue, which is why they must be forced to stop. Those who are negatively affected by the pollution are an important constituency when it comes to generating the political will necessary to impose these sanctions. This is true in the case of climate change, but the difficulty is compounded by the fact that it is in interest, narrowly construed, of most people alive today to continue producing carbon pollution. Furthermore, it is likely to be in the interest of those alive in 10 years, 50 years or 100 years, to continue polluting. Each generation finds itself in a situation where the die has already been cast, as far as its own quality of life is concerned, and so its choices affect primarily those who have not yet been born. As a result, the most important constituency to benefit from mitigation cannot be mobilized. We must therefore force ourselves to stop polluting, even though few of us have any incentive to do so. It is difficult to think of many other policy domains that rely so completely on rational insight, rather than interest group mobilization, to motivate implementation of a solution. The practical challenge is clear.

What I have attempted to show in this paper is that the intergenerational structure of the problem creates not only practical difficulties but theoretical ones as well. The major challenges cannot be addressed through an "off the shelf" application of the most widely held normative theories. Perhaps Parfit's most useful contribution in this domain is to have shown how fundamentally strange certain aspects of our moral relations to future people are. The most common mistake made by subsequent philosophers has been to underestimate how great an impact these unusual features of our relations to future people have on the climate change issue. The various attempts to extend contemporaneous moral theories to accommodate future people or their interests illustrate the problems that stem from such an underestimation. Climate ethicists have therefore struggled mightily to find a normative framework that will avoid generating absurd or unnecessarily controversial conclusions. This has resulted in a funneling of the discussion toward a small set of views that retain some plausibility. My goal in this paper has been to provide an overview of the considerations that have been driving this discussion, in part to encourage those doing constructive work in this domain to focus their energies on the development of normative approaches that have some chance of success.

Joseph Heath is Professor in the Department of Philosophy at the University of Toronto. He is the author of several books, including Philosophical Foundations of Climate Change Policy (Oxford, 2021).

References

Adler, M. D., Anthoff, D., Bosetti, V., Garner, G., Keller, K., & Treich, N. (2017). Priority for the worse-off and the social cost of carbon. *Nature Climate Change*, 7(6), 443–449. https://doi.org/10.1038/nclimate3298

Adler, M. D., & Treich, N. (2015). Prioritarianism and climate change. *Environmental and Resource Economics*, 62(2), 279–308. https://doi.org/10.1007/s10640-015-9960-7

Arntzenius, F., Elga, A., & Hawthorne, J. (2004). Bayesianism, infinite decisions, and binding. *Mind*, 113(450), 251–283. https://doi.org/10.1093/mind/113.450.251

Arrow, K. (1999). Discounting, morality, and gaming. In P. R. Portney & J. P. Weyant (Eds.), *Discounting and intergenerational equity* (pp. 13–21). Resources for the Future.

Barry, B. (1999). Sustainability and intergenerational justice. In A. Dobson (Ed.), *Fairness and futurity* (pp. 93–117). Oxford University Press. https://doi.org/10.1093/0198294891.003.0005

Berger, J. (2001, March 29). Bush kills climate treaty. *The Guardian*.
Bilal, A., & Känzig, D. R. (2026). The macroeconomic impact of climate change: Global versus local temperature. *The Quarterly Journal of Economics*, 141(2), 889–944. https://doi.org/10.1093/qje/qjag011
Boran, I. (2018). *Political theory and global climate action*. Routledge.
Boran, I. (2019). On inquiry into climate justice. In T. Jafry (Ed.), *Routledge handbook of climate justice* (pp. 26–41). Routledge.
Bradsher, K. (2024, March 7). How China came to dominate the world in solar energy. *New York Times*.
Broome, J. (1992). *Counting the costs of global warming*. White Horse Press.
Broome, J. (1994). Discounting the future. *Philosophy and Public Affairs*, 23(2), 128–156. https://doi.org/10.1111/j.1088-4963.1994.tb00008.x
Broome, J. (2013). *Climate matters*. W. W. Norton.
Brundtland, G. H. (1987). *Our common future*. United Nations.
Buchholz, W., & Schumacher, J. (2010). Discounting and welfare analysis over time: Choosing the η. *European Journal of Political Economy*, 26(3), 372–385. https://doi.org/10.1016/j.ejpoleco.2009.11.011
Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577), 235–239. https://doi.org/10.1038/nature15725
Bykvist, K. 2007. Preference formation and intergenerational justice. In A. Gosseries, & L.H. Meyer (Eds.), *Intergenerational justice*. Oxford University Press.
Caney, S. (2010). Climate change, human rights, and moral thresholds. In S. Gardiner, S. Caney, D. Jamieson, & H. Shue (Eds.), *Climate ethics* (pp. 163–177). Oxford University Press.
Carroll, L. (1981). *Alice's adventures in Wonderland & through the looking glass*. Bantam Classic.
Chen, L., Feng, P., Shaver, L., & Wang, Z. (2021). Maternal mortality ratio in China from 1990 to 2019: Trends, causes and correlations. *BMC Public Health*, 21(1), 1536. https://doi.org/10.1186/s12889-021-11557-3
Coase, R. (1960). The problem of social cost. *Journal of Law and Economics*, 3, 1–44. https://doi.org/10.1086/466560
Cowen, T. (1992). Consequentialism implies a zero rate of discount. In P. Laslett & J. S. Fishkin (Eds.), *Justice between age groups and generations* (pp. 162–168). Yale University Press.
Crippa, M., Guizzardi, D., Pagani, F., & Banja, M. (2024). *GHG emissions of all world countries*. Publications Office of the European Union.
Dasgupta, P. (2008). Discounting climate change. *Journal of Risk and Uncertainty*, 37(2-3), 141–169. https://doi.org/10.1007/s11166-008-9049-6
Davidson, M. (2008). Wrongful harm to future generations: The case of climate change. *Environmental Values*, 17(4), 471–488. https://doi.org/10.3197/096327108X368494
DesRoches, C. T. (2019). On the concept and conservation of critical natural capital. *International Studies in the Philosophy of Science*, 32(3–4), 207–228. https://doi.org/10.1080/02698595.2020.1788347
Ferranna, M. (2021). Discounting under risk: utilitarianism vs. prioritarianism. In M. Budolfson, T. McPherson, & D. Plunkett (Eds.), *Philosophy and climate change* (pp. 67–89). Oxford University Press.
Gardiner, S. M. (2011). *A perfect moral storm*. Oxford University Press.
Gaspart, F., & Gosseries, A. (2007). Are generational savings unjust? *Politics, Philosophy and Economics*, 6(2), 193–217. https://doi.org/10.1177/1470594X07073006
Gollier, C. (2006). *An evaluation of Stern's report on the economics of climate change.* (IDEI Working Paper 464, pp. 1–10).
Gosseries, A., & Meyer, L. H. (2009). *Intergenerational justice*. Oxford University Press.
Groom, B., & Maddison Pr, D. (2019). New estimates of the elasticity of marginal utility for the UK. *Environmental and Resource Economics*, 72(4), 1155–1182. https://doi.org/10.1007/s10640-018-0242-z
Hamilton, C. (2011). *Requiem for a species*. Routledge.
Heath, J. (2021). *Philosophical foundations of climate change policy*. Oxford University Press.
Jamieson, D. (2014). *Reason in a dark time*. Oxford University Press.
Kagan, S. (1984). Does consequentialism demand too much? Recent work on the limits of obligation. *Philosophy and Public Affairs*, 13(3), 239–254.
Karnein, A. (2022). Rawls and the future: On the possibility of cooperation across time. *Philosophy and Public Affairs*, 50(3), 271–300. https://doi.org/10.1111/papa.12213
Kelleher, J. P. (2025). *The social cost of carbon*. Oxford University Press.
Keynes, J. M. (2010). *Essays in persuasion*. Palgrave Macmillan.
Koopmans, T. (1960). Stationary ordinal utility and impatience. *Econometrica*, 28(2), 287–309. https://doi.org/10.2307/1907722
MacAskill, W. (2022). *What we owe the future*. Hachette.
Mintz-Woo, K. (2019). Principled utility discounting under risk. *Moral Philosophy and Policy*, 6(1), 89–112. https://doi.org/10.1515/mopp-2018-0060
Moellendorf, D. (2014). *The moral challenge of dangerous climate change*. Cambridge University Press.
Moellendorf, D. (2022). *Mobilizing hope*. Oxford University Press.
Mulgan, T. (2011). *Ethics for a broken world*. McGill-Queen's University Press.

Neumayer, E. (2010). *Weak versus strong sustainability* (3rd ed.). Edward Elgar.
Nordhaus, W. D. (2007). A review of the Stern review on the economics of climate change. *Journal of Economic Literature*, 45(3), 686–702. https://doi.org/10.1257/jel.45.3.686
Norton, B. G. (2005). *Sustainability*. University of Chicago Press.
Nozick, R. (1971). *Anarchy, state, and utopia*. Basic Books.
O'Neill, B. C. (2023). Envisioning a future with climate change. *Nature Climate Change*, 13(9), 874–876. https://doi.org/10.1038/s41558-023-01784-4
Parfit, D. (1982). Future generations: Further problems. *Philosophy and Public Affairs*, 11(2), 113–172.
Parfit, D. (1984). *Reasons and persons*. Oxford University Press.
Parfit, D. (2002). Equality or priority? In M. Clayton & A. Williams (Eds.), *The ideal of equality* (pp. 81–125). Palgrave.
Parfit, D. (2017). Future people, the non-identity problem, and person-affecting principles. *Philosophy and Public Affairs*, 45(2), 118–157. https://doi.org/10.1111/papa.12088
Pollock, J. L. (1983). How do you maximize expectation value? *Noûs*, 17(3), 409–421. https://doi.org/10.2307/2215257
Posner, E. A., & Weisbach, D. (2010). *Climate change justice*. Princeton University Press.
Rawls, J. (1971). *A theory of justice*. Harvard University Press.
Rawls, J. (2001). *Justice as fairness*. Harvard University Press.
Rendall, M. (2018). Discounting and the paradox of the indefinitely postponed splurge. In R. Kanbur, & H. Shue (Eds.), *Climate justice: Integrating economics and philosophy* (pp. 175–187). Oxford University Press.
Ritchie, H. (2024). *Not the end of the world*. Hachette.
Shue, H. (1999). Bequeathing hazards: Security rights and property rights of future humans. In M. Dore, & T. Mount (Eds.), *Global environmental economics: Equity and limits to markets* (pp. 38–52). Blackwell.
Sidgwick, H. (1981). *The methods of ethics* (7th ed.). Hackett.
Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, 1(3), 229–243.
Smart, J. J. C., & Williams, B. (2013). *Utilitarianism: For and against*. Cambridge University Press.
Solow, R. (1974). Intergenerational equity and exhaustible resources. *Review of Economic Studies*, 41(5), 29–45. https://doi.org/10.2307/2296370
Stanczyk, L. (2021). How quickly should the world reduce its greenhouse gas emissions? Climate change and the structure of intergenerational justice. In M. Budolfson, T. McPherson, & D. Plunkett (Eds.), *Philosophy and climate change* (pp. 295–322). Oxford University Press.
Stern, N. (2014). *The economics of climate change: The Stern review*. Cambridge University Press.
Táíwò, O. O. (2022). *Reconsidering reparations*. Oxford University Press.
Tarsney, C. (2017). Does a discount rate measure the costs of climate change? *Economics and Philosophy*, 33(3), 337–365. https://doi.org/10.1017/S0266267117000049
von Allmen, K. (2022). Reducing climate change harms: How to make remedial responsibilities applicable. *Public Affairs Quarterly*, 36(4), 325–352. https://doi.org/10.5406/21520542.36.4.04
Waidelich, P., Batibeniz, F., Rising, J., Kikstra, J. S., & Seneviratne, S. I. (2024). Climate damage projections beyond annual temperature. *Nature Climate Change (London)*, 14(6), 592–599. https://doi.org/10.1038/s41558-024-01990-8
Wallace-Wells, D. (2018). *The uninhabitable earth: Life after warming*. Columbia University Press.
World Bank. (2022). *Four decades of poverty reduction in China*. The World Bank.
Zwarthoed, D. (2015). Cheap preferences and intergenerational justice. *Revue de philosophie économique*, 16(1), 69–101. https://doi.org/10.3917/rpec.161.0069

Cite this article: Heath, J. 2026. How Climate Change Confounds Conventional Normative Theories. *Canadian Journal of Philosophy* 56: e21, 1–19, doi:10.1017/can.2026.10113