Simon Dietz, Ben Groom, and William A. Pizer

Weighing the costs and benefits of climate change to our children

Article (Published version) (Refereed)

Weighing the Costs and Benefits of Climate Change to Our Children

Simon Dietz, Ben Groom, and William A. Pizer

Summary

Our efforts to put the brakes on climate change or adapt to a warming climate present a fundamental tradeoff between costs borne today and benefits that accrue to the children and grandchildren of the current generation. In making investments today that affect future generations' prospects, we need to think carefully about how we value their welfare compared to our own.

A common economic formula recommends giving up only 5 cents today for every dollar of benefits 100 years in the future; we call this *discounting* the future. Underlying this approach is the assumption that future generations will be much better off than our own, just as we are much wealthier than our ancestors were. Would our descendants' agree with this approach? Are there reasons to put more value on future benefits?

William Pizer, Ben Groom, and Simon Dietz discuss three possible reasons that we might put a higher value on future benefits. First, people disagree considerably about the correct discount rate. Other plausible interpretations of society's preferences or observed data could increase the weight we place on future benefits by as much as a factor of five. Second, we may have failed to correctly value future climate change impacts, particularly those related to the loss of environmental amenities that have no close monetary substitutes. Third, we may not be properly valuing the risk that a warming climate could cause sudden and catastrophic changes that would drastically alter the size of the population.

Ultimately, the authors write, many of the choices about how we value future generations' welfare come down to ethical questions, and many of the decisions we must make come down to societal preferences—all of which will be difficult to extract from data or theory.

www.futureofchildren.org

Simon Dietz is a professor in the Department of Geography and Environment, director of the ESRC Centre for Climate Change Economics and Policy, and co-director of the Grantham Research Institute on Climate Change and the Environment at the London School of Economics and Political Science. Ben Groom is an associate professor in the Department of Geography and Environment at the London School of Economics and Political Science. William A. Pizer is a professor in the Sanford School of Public Policy and a faculty fellow in the Nicholas Institute for Environmental Policy Solutions at Duke University.

Charles Kolstad of Stanford University reviewed and critiqued a draft of this article.

Future generations are the current generation's children, grandchildren, and so on. That intergenerational perspective gives rise to extremely thorny questions about how to evaluate and make trade-offs between the wellbeing of current generations and the wellbeing of their descendants. In the context of climate change, we can't avoid intergenerational comparisons because greenhouse gas emissions today produce impacts that will last for hundreds of years. Therefore, we must analyze trade-offs over extremely long time horizons. In short, the payoffs from our own costly mitigation efforts will accrue to our children and their descendants. As will be made plain, to make decisions in the face of such dynamics, we must carefully analyze efficiency and equity. In particular, we can imagine asking whether our children will look back and take issue with how we valued their welfare compared with our own.

The principal economic tool for decision making is cost–benefit analysis (CBA). In a CBA, all current and future costs and benefits, or net benefits, in each period are given a weight and are then summed, with costs entered as negative benefits. Policy options with higher net benefits are generally preferred. Current costs and benefits have a weight of one. The weight placed on future costs and benefits is determined by a number known as the discount rate and, more specifically, in the case of societal rather than private decision making, the social discount rate (SDR). The SDR determines how quickly the weight placed on future costs and benefits diminishes with the time horizon being considered: the higher the SDR, the lower the influence of future costs and benefits on present values. When we consider long time horizons, as we must with climate change, small changes in the SDR can lead to extremely large differences in the weight we place on future costs and benefits.

What determines the SDR depends on how we conceive of social welfare across time and/or generations in the first place. The standard CBA approach is grounded in a welfare framework known as discounted utilitarianism (DU). In DU, welfare in future years and for future generations is added together, with future generations effectively viewed as extensions of ourselves further into the future (a *representative-agent* approach). A key feature of DU is that we value additional dollars less as we become richer. Coupled with the assumption that continued economic growth will leave our future selves (children and grandchildren) better off, this tends to support significant discounting of dollar-valued benefits in the future. But how much discounting?

Small changes in the social discount rate can lead to extremely large differences in the weight we place on future costs and benefits.

As we'll see, a typical application of the DU framework leads to weighting dollar-valued costs and benefits a hundred years in the future at roughly one-twentieth of the value of similar costs and benefits today. Some people, perhaps even our descendants themselves, might view such a weighting as incorrect or inequitable. Therefore, we explore a variety of reasons we might make the weighting more favorable for future generations.

Most significantly, scholars who work on DU disagree about which parameters we should use to reach the SDR and how those parameters should be calculated. One disagreement has been over whether to use a normative approach or a positive approach. Broadly speaking, the normative school asks the ethical question, "How *ought* we trade off our own welfare with that of our descendants?" The positive school asks instead, "Empirically, how *do* we trade off current and future welfare?" Within both schools, scholars further disagree about how to interpret the evidence and apply the ethical judgments that determine the SDR.

Beyond the question of parameters, a number of extensions and alternative conceptions of social welfare across time (that is, intertemporal social welfare) can affect future valuations. One key extension explicitly considers the significant uncertainty around future economic growth and welfare—in our case, economic growth and welfare distinct from climate change's effects. Another extension considers whether environmental resources can be substituted. We can also abandon the DU model altogether and consider other ways to assess social welfare across time and generations that are rooted in alternative conceptions of fairness and justice.

We could also imagine that the effects of climate change on human health and mortality could be so serious as to affect the size of the population, meaning that our choices about climate mitigation would affect not only how well off our descendants would be but also how many of them there would be. That possibility raises yet other ethical issues, such as whether it's better to have large, subsistence-level populations or small, better-off ones. Thinking through such issues is a difficult task that has enormous consequences for the weight we place on our descendants when we evaluate intergenerational policies like climate mitigation.

The Intergenerational Trade-Off

The slow pace, or inertia, of the climate system's response to greenhouse gas emissions implies that an intergenerational trade-off lies at the heart of questions about how much to cut those emissions. (See the article in this issue by Michael Oppenheimer and Jesse Antilla-Hughes for more about the science of climate change.) Arguably, no research has presented the trade-off issue more starkly than the set of economic models built to simultaneously investigate the costs and benefits of reducing emissions and how those costs and benefits are distributed across time. We'll take the best-known model—the DICE (Dynamic Integrated Climate-Economy) model built by Yale economist William Nordhaus—and illustrate the important points by way of a few scenarios.[1]

Figure 1 plots the baseline level of aggregate consumption—essentially, how wealthy the world becomes over time in the absence of any action. Focusing on the "standard" case, we see the typical assumption that, over time, the world becomes much, much richer.

Figure 2 plots the net benefits of a particular mitigation scenario (relative to inaction, or, in other words, the baseline) over the next two centuries, according to DICE. The vertical axis shows the net benefits in each period as a share of consumption (top panel) and as valued in trillions of dollars (bottom panel) in that period; negative numbers indicate net

Figure 1. Aggregate Consumption under Standard and Catastrophic Climate Change Scenarios

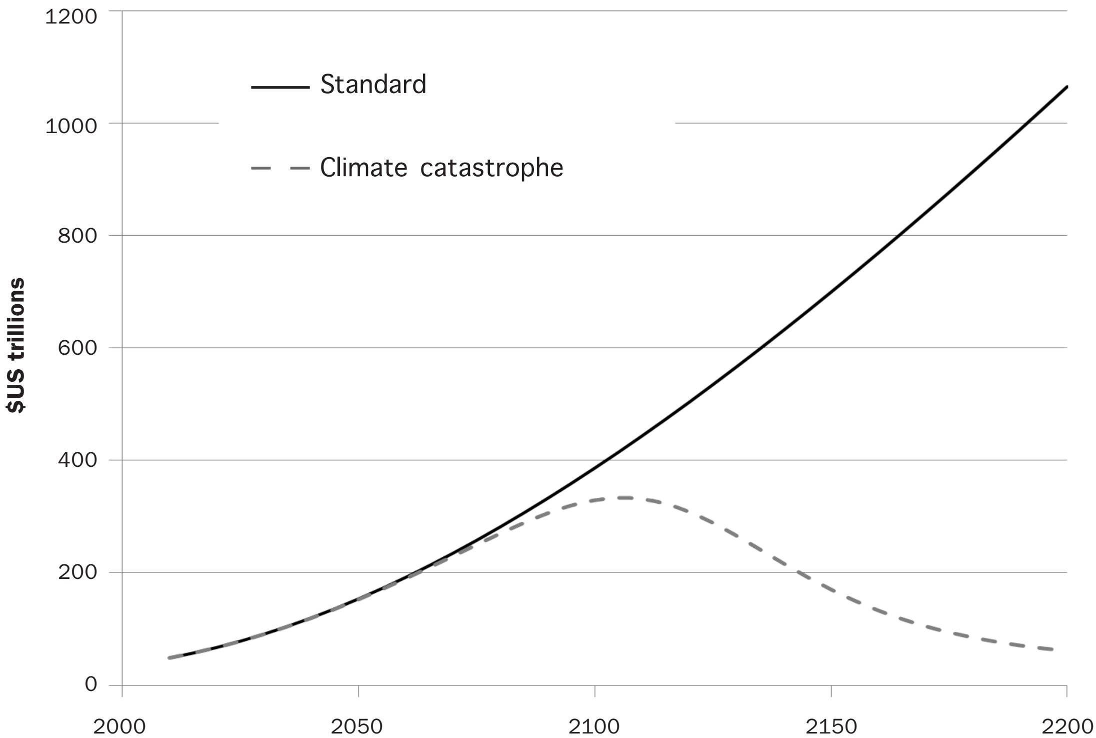

Source: A version of the DICE model built by Dietz et al., which extends Nordhaus's DICE-2013R model. See Simon Dietz, Christian Gollier and Louise Kessler, "The Climate Beta," Working Paper no. 215 [Centre for Climate Change Economics and Policy] and no. 190 [Grantham Research Institute on Climate Change and the Environment] (London School of Economics and Political Science, London, UK, 2015) and William D. Nordhaus, "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches," *Journal of the Association of Environmental and Resource Economists* 1 (2014): 273–312.

costs in that period. This particular mitigation scenario is intended to hold the increase in global mean temperature below 2º Celsius (3.6º Fahrenheit). The two different lines reflect different sets of assumptions about the costs and benefits of such a scenario.[2]

The "standard" case in figure 2 broadly reflects Nordhaus's usual parameter assumptions, which are typical of most research on the topic—at least until recently. The intergenerational trade-off in this case becomes immediately clear: for the rest of this century, society will have to sacrifice income—up to nearly 4 percent of baseline consumption in 2060—to avoid damages and adaptation costs from climate change, which occur mostly after 2100. Absent any weighting to reflect time preferences, the cumulative net benefits of mitigation are larger than the net costs. That is, in the standard case, the positive area under the curve after 2100 exceeds the negative area above the curve before 2100. However, applying the sort of SDR that governments routinely use would substantially reduce the present value of long-term benefits compared with the near-term costs. The SDR that balances the present value of costs and benefits in this scenario—referred to as the internal rate of return—is 2.9 percent, which is slightly lower than the 3 percent rate that the US government used in a recent

Figure 2. Net Benefits of a Global Emissions Path to Avoid More than 2°C Warming

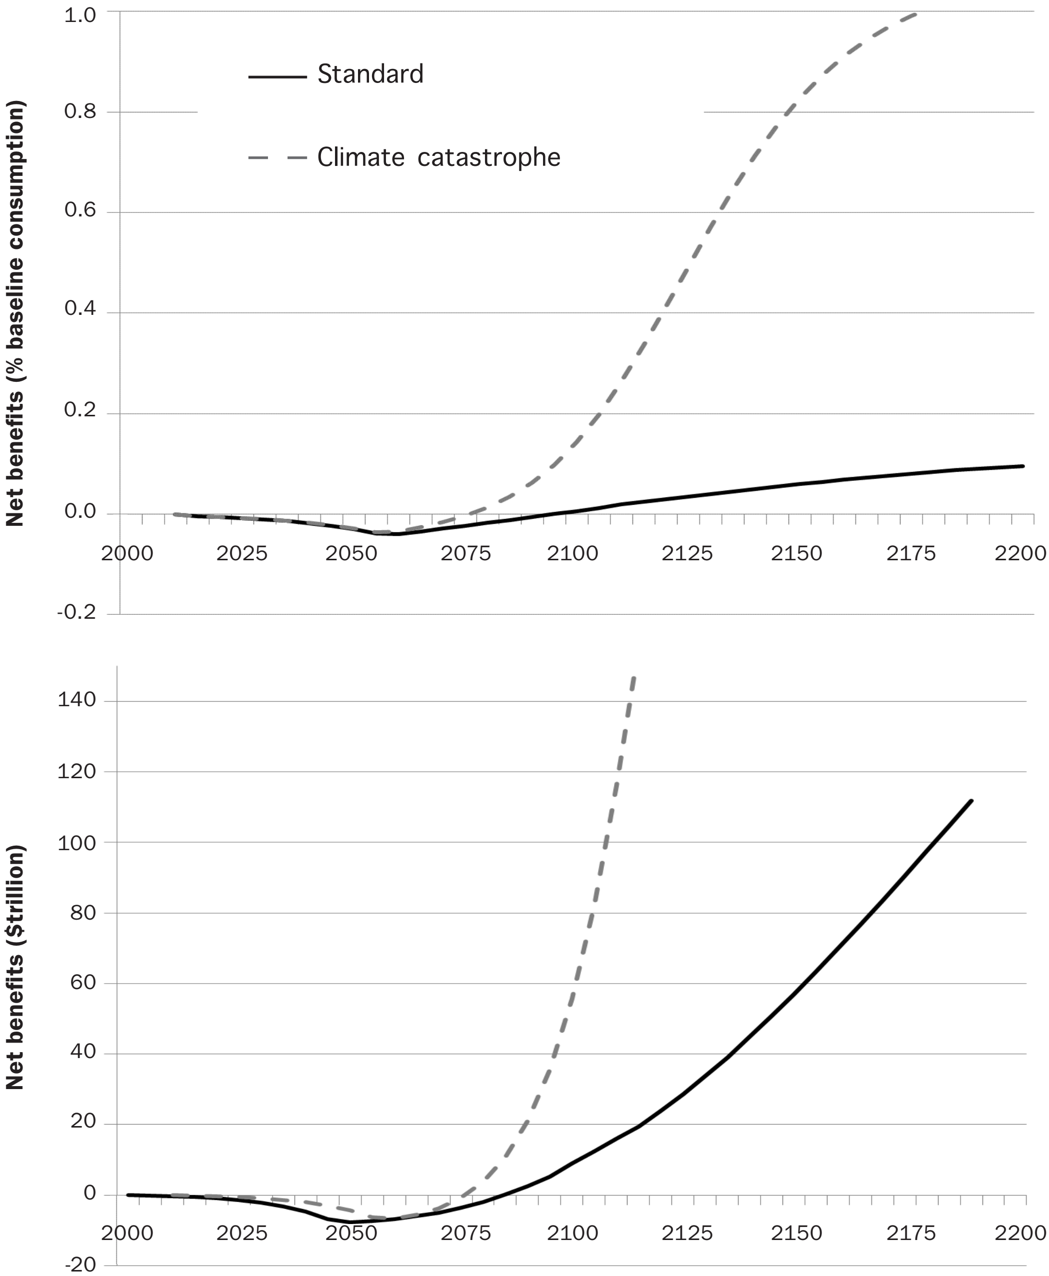

Source: A version of the DICE model built by Dietz et al., which extends Nordhaus's DICE-2013R model. See Simon Dietz, Christian Gollier and Louise Kessler, "The Climate Beta," Working Paper no. 215 [Centre for Climate Change Economics and Policy] and no. 190 [Grantham Research Institute on Climate Change and the Environment] (London School of Economics and Political Science, London, UK, 2015) and William D. Nordhaus, "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches," *Journal of the Association of Environmental and Resource Economists* 1 (2014): 273–312.

analysis of climate change benefits and much lower than the 4 to 5 percent assumed by Nordhaus himself.[3] In other words, under standard assumptions, the rate of return to a societal investment that would keep global mean temperature change below 2ºC is less than what is required by typical SDR values. A global effort to reduce emissions by that much would fail a cost–benefit test, meaning that it would not increase social welfare.

Let's now explore how robust that result is. The "climate catastrophe" scenario in figures 1 and 2 combines a very sensitive response by the global mean temperature to rising atmospheric greenhouse gas concentrations with damages that rise particularly steeply in response to warming—a double whammy, if you like. As a result, human wellbeing initially grows but then falls back to current levels by the end of 2200. That situation contrasts sharply with the standard assumptions, wherein the economy grows by a factor of more than 10, with or without climate change (see figure 1). In the climate catastrophe scenario, the net benefits of mitigation that limits climate change to 2ºC skyrocket (see figure 2); they are already more than 10 percent of consumption by 2100. In that case, not only does it make economic sense to keep global warming under 2ºC; it makes sense to do a whole lot more.

Herein lies our main observation from the DICE model: the goal of stabilizing temperatures at 2ºC doesn't pass a cost–benefit analysis that uses standard assumptions about climate change and an SDR of 3 percent, the rate suggested by the US government. A 3 percent rate implies that dollar impacts in 2115 are weighted at roughly one-twentieth the value of dollar impacts in 2015. Yet even with that kind of weighting, assumptions of direr climate change consequences can overturn that result.

In this article, most of our interest centers on whether weighting consequences a hundred years in the future at roughly one-twentieth the value of consequences today is correct and/or fair. But before tackling that issue, we briefly step beyond DICE to ask, "What do we know about the costs of mitigation, as well as climate change's impacts and the benefits of mitigation?"

The Costs of Mitigation

What does it mean to say that mitigating climate change is costly? Ultimately, households and their children pay for mitigation through reductions in the welfare they enjoy. Policies to reduce emissions will raise the price of producing energy as we seek less-carbon-intensive alternatives. That means households will pay more not only for the energy they use but also for energy-intensive goods such as products made of steel and aluminum. In response, companies and households will switch from carbon-intensive, or dirty, production technologies to low-carbon, or clean, counterparts and/or look for ways to simply use less electricity. However, all of those things have welfare costs. The alternatives either cost more or entail loss of service or convenience. Though there may be negative-cost opportunities that can reduce emissions and raise welfare at the same time, most economists are skeptical that such opportunities will be plentiful or easy to capture.[4]

It's also tempting to imagine that costs to businesses somehow won't affect households. However, business losses come back to households in several ways, including lower income from stocks, reduced wages, or increased prices that work their way through the whole economy. For example, if the price of electricity goes up, so do the prices of aluminum and, in turn, aluminum products such as foil and cans.

It's helpful to understand those basic linkages not only to make the notion of mitigation costs tangible but also because the scholarship on mitigation costs follows various approaches and arrives at various ways to measure cost, not all of which are easily comparable. For instance, one popular way to roughly approximate mitigation costs is to (1) estimate the cost and potential to reduce emissions of each of a menu of technologies and options (for example, energy-efficient refrigerators or wind power), (2) sort them from least to most expensive, and (3) add them up until the desired level of abatement is reached.[5] That approach has the advantage of extensive detail about low-carbon technologies. At the same time, since all of those options are being estimated separately, they might not add up as a whole, because implementing some strategies might affect the cost of others. Perhaps the best example is how the emission-reducing benefits of using less electricity would decline if we switched to less-emission-intensive electricity sources.

By contrast, aggregate economic models emphasize relationships between energy prices and the supply and demand (from producers and consumers, respectively) for energy products. By integrating the behavior of such actors, an aggregate model ensures that everything adds up—but at the expense of retaining little detail about the various technologies for cutting emissions.

Stylized Facts

The extensive research on mitigation costs is dominated by simulation results from integrated economy-energy-climate models. *The Fifth Assessment Report* by the Intergovernmental Panel on Climate Change (IPCC) does a very good job of summarizing the vast majority this research.

- Most models predict that climate mitigation is costly and would reduce economic growth when considered against a "utopian" scenario wherein no mitigation takes place and climate change has no effect, but the reduction in growth prospects tends to be relatively small, except perhaps for the most stringent mitigation targets.[6] By 2050, the loss in consumption relative to a utopian scenario ranges from about 0.5 to just over 5 percent, depending on the depth of emissions cuts.
- Although it's useful to think of mitigation as a simple investment whereby we spend a onetime sum in the beginning to receive a stream of benefits in the future, mitigation is in fact an ongoing activity, and its global costs will rise over time. In this way, our children and their children will each face the dilemma of how to weight the costs they will endure to benefit future generations, even as they reap whatever benefits accrue from our own efforts.
- The global costs of mitigation increase with the stringency of the emissions target. Some evidence suggests that they increase more than proportionately.[7]
- Models' forecasts differ widely in what it would cost to achieve the same emissions target. And the differences increase as levels of required emissions reductions increase. That divergence has many causes, including different assumptions about population growth and economic

growth, but a particularly important set of assumptions concerns the availability and costs of low-carbon technologies.[8]

- If key mitigation technologies aren't available at a reasonable cost (or aren't available at all), global mitigation costs could increase significantly. Carbon capture and storage (CCS) is a notable example. CCS is a technology for capturing the carbon dioxide from large emitting facilities and storing it in underground geologic formations. It makes continued fossil-fuel burning consistent with emissions targets in the short term; it can be combined with a range of emissions-generating technologies (that is, not just electricity generation); and it can even be combined with biofuels to yield negative emissions. But as of today, just 13 CCS facilities are operating in the world, and only one is attached to a power plant.[9] When we purposely eliminate CCS from a range of models, they predict that mitigation will cost much more.

- Delaying efforts to reduce emissions in the coming years will increase mitigation costs further in the future, partly because we'll have to make deeper cuts later on and partly because we'll have locked in carbon-intensive infrastructure in the intervening period. Many emissions are caused by very long-term investments that are in turn tied to very long-term infrastructure. For example, cars might have a useful life of 10 or 15 years, but the fueling infrastructure is more durable. If the task is simply to hit a given emissions target at the lowest cost, too much delay is, consequently, a bad thing.

- Global mitigation costs increase if some countries don't pull their weight. In part, that's simply because some mitigation that could have taken place cost-effectively in those countries will now have to take place at a higher cost elsewhere. But it also reflects the phenomenon of carbon leakage: countries that lag in restricting emissions may attract carbon-intensive industries, thereby increasing aggregate emissions and making other countries work that much harder.[10]

- The costs of mitigation will vary from country to country, for two reasons. First, opportunities vary. Some countries are blessed with renewable resources, and others are not. Some countries are already building new infrastructure, and others would have to retire existing facilities before the end of their useful lives. That points to significant mitigation potential in large and fast-growing developing countries such as China and India. Second, national governments—both alone and in various multilateral settings—will make decisions about how to financially support efforts in poorer countries. Ultimately, the distribution of costs across countries will depend on some combination of where the mitigation opportunities exist, how much those countries are willing to spend to mitigate, and how much money other countries are willing to provide to meet the costs of mitigation.

- Estimates from integrated models don't include all of the factors thought to affect mitigation costs. Most models assume that, apart from climate change itself, the economy is otherwise functioning perfectly, or, if it isn't, the number of ways it could malfunction is at least severely limited. Fully incorporating imperfections into the functioning of the economy—something we can only imagine being able to do—could either increase or decrease global mitigation costs. One example is the cobenefits that mitigation would have for public health via reduced emissions of conventional air pollutants; in some parts of the world, those cobenefits could be very substantial, thus decreasing the cost of mitigation.[11] (For more about the relationship between emissions and health, see the article by Allison Larr and Matthew Neidell elsewhere in this issue.) On the other hand, economists have argued that carbon regulation would increase the cost of the existing tax system, thereby adding to the cost of mitigation.[12]

Costs of Climate Change

The costs of climate change fall into two categories: the costs of adapting to climate change (for instance, by increasing defenses against coastal flooding) and the costs of residual damage from climate change after adaptation, such as flooding from a storm surge that overtops those strengthened coastal defenses.

Costs can be further subdivided in a number of ways, two of which are helpful for understanding the nature of climate impacts. The first is simply to categorize costs by the sector of the economy in which they fall, which quickly leads to the conclusion that a few sectors, such as agriculture and forestry, are especially vulnerable to gradual climate change. However, increases both in weather's variability and in instances of extreme weather have the potential to affect a wider range of activities. For instance, Japanese automobile manufacturer Toyota suffered disruption to its supply chain when Bangkok was flooded in 2011.

The second subdivision is between so-called market costs and nonmarket costs. Market costs are costs paid in the real economy, such as losses in agricultural output and increased expenditure on air conditioning. Nonmarket costs are impacts that are real but nonetheless aren't paid directly in the real economy—notably, the value most people would put on lost human health and damage to the natural environment beyond simple market losses. Not all research explicitly distinguishes between the two, but some important work has shown that nonmarket impacts, when rendered equivalent to market impacts by a technique called *shadow pricing*, are relatively substantial—perhaps greater than market costs.[13] We return to this topic at the end of the next section.

Stylized Facts

Research on the impacts of climate change is voluminous and, partly because those impacts are so diverse, much of it focuses on a particular impact. For example, one researcher might build a crop model to analyze agricultural impacts. By contrast, few economic models seek to aggregate impacts. The research as a whole was recently summarized in Working Group II's contribution to the IPCC's *Fifth Assessment Report*.

- The first few degrees of warming will bring costs to some and benefits to others (for example, increased agricultural productivity at high latitudes in the Northern Hemisphere). Consequently, models disagree on whether the initial global costs of climate change are positive or negative overall, though most find that they're negative. At warming of 2 or 3ºC above the preindustrial level, which for many studies has become a benchmark, the small set of integrated models estimates a global cost in the range of minus 0.1 to 3 percent of gross domestic product relative to no climate change.[14]

- Tropical and subtropical developing countries are relatively more vulnerable to climate change. They're more exposed to adverse changes in climatic conditions; they're more sensitive to climate change because a larger share of their economic activity takes place in climate-sensitive sectors, particularly agriculture; and they have less capacity to adapt.[15]

- Beyond three degrees of warming, we understand little about the impacts of climate change, particularly at the aggregate level. Some integrated economic models continue to estimate small costs relative to the global economy, but others predict spiraling costs that would eventually lead to a global economic catastrophe.[16] The models agree, however, that losses accelerate as warming increases.

- No model covers all of the known effects of climate change. We can merely speculate on what would happen if we included those omitted variables, though most scholars agree that they would increase costs, because the omitted effects include some of the most worrying ones, such as climate-induced conflict.[17]

All of those observations about costs and benefits suggest that we should approach any estimates with caution. Nonetheless, most stakeholders and governments see value in trying to predict climate change's effects, even if the estimates are flawed. With a sketch in hand of such calculations concerning today's costs and future benefits, we now turn to the question of impacts across time and generations. On what basis do we judge changes in our wellbeing versus the wellbeing of our descendants and their children?

Evaluating Our Descendants' Wellbeing

To decide whether investing for the future is somehow "better" or "worse" for society, we need ways to evaluate *better* and *worse*. It boils down to a definition of societal wellbeing across time and multiple generations against which to compare different courses of action. Only then can we evaluate whether the costs of mitigating climate change outlined in the previous section would be outweighed by the benefits to our descendants and whether the investment in future generations is "worth it."

Discounted Utilitarianism

The standard approach in cost–benefit analysis is to weight costs and benefits at different points in time using an SDR. But lurking behind the definition of an SDR is a larger notion of welfare, which recognizes that changes in income or dollars alone may not be the best way to measure the degree to which a particular person is better off. Instead, welfare economics defines individual wellbeing in terms of utility. The basic difference that such a definition of wellbeing introduces is that, although income and consumption contribute to wellbeing, extra income's contribution to wellbeing diminishes as an individual or society gets richer. Utility can also capture the idea that wellbeing may depend on things other than the market goods that we consume, such as clean air and good health.

But how should we aggregate welfare for many people—particularly those living at different points in time? That is, how do we add up and then compare the welfare effects of courses of action that may affect both today's society and future societies populated by our descendants? There are many ways we might do this, but the standard approach is to define the utility of an "average" person over a period of time—say, a year—in terms of consumption of market goods and, possibly, nonmarket goods such as environment and health. Those annual utilities of an average—or *representative*—agent are added together to obtain an overall measure of intertemporal social welfare (a number). Each generation and each person is assumed to have the same utility function. Thus, average utility is typically multiplied by the number of people alive in each period before current utility and future utility are added together. Future utility may or may not be discounted. This additive, representative-agent approach is discounted utilitarianism, and it is the standard approach involved in the welfare economics underpinning CBA and the economics of climate change.

Within the DU approach lie two essential issues that determine how much weight to place on the monetized costs and benefits accruing to future generations versus our own. The first issue is that society may place different weights on utility in future years. That's relatively uncontroversial when applied to an individual, because most people prefer to receive a net benefit earlier, all else equal. But when we extend the principle to different generations, questions of equity arise. For example, imagine that two generations—our own today and one in the future—enjoy the same level of income and hence of utility (in this standard approach). Moreover, a particular monetary benefit would lead to the same increase in utility for both. From today's perspective, when aggregating and adding those utilities together to measure social welfare, we might wish to place less weight on the future generation's utility than on our own today. That is, all else equal, society might prefer a given monetary benefit if it is delivered to this generation rather than the next.

Two essential issues determine how much weight to place on the costs and benefits accruing to future generations. The first is that society may place different weights on utility in future years. The second concerns aversion to income inequality.

Alternatively, from the perspective of equal treatment of generations, we might not want to discount the future generation's utility at all when making the welfare calculation. This is the first essential issue of intergenerational equity that we must face when deciding how to evaluate the costs and benefits of climate mitigation, and it has caused a great deal of debate within the utilitarian tradition and beyond. We return to this debate later.

The second issue concerns aversion to income inequality. From a societal perspective, a given addition to income for a poor person is typically thought to raise welfare more than the same addition of income would for a rich person. An intervention (for example, a public infrastructure project or a climate change mitigation project) that yielded incremental income to the poor would then be worth more to society than would a project that yielded the same incremental income only to the rich. For example, consider an intervention that lowers present income by $1 in 2015 and raises future income by $1 in 2115. Let's say that the present society has an income of $10,000 per capita, and that the future society of our descendants has an income of $30,000 per capita. That is, income has grown over time, and the future society is richer as a consequence. In our example there is no inflation, so the income growth is real; for simplicity, imagine that we are considering only a single person at each point in time. Here, the considered intervention would lower welfare. Why? Because $1 in the future is worth less than $1 today, solely because the future is wealthier (has higher income) and we are averse to income inequality. This is the second reason we might wish to discount future costs and benefits. Of course, that wealth effect is a double-edged sword. If the future generation is poorer—that is, if growth is negative, as has been the case in many developing countries over the past 30 years—then $1 in the future contributes more to social welfare than $1 today does.

Typically, those two reasons for putting less weight on future generations combine to form an SDR that indicates the rate at which the weight we place on future generations' consumption declines the further we look into the future. We call it a social discount rate because the context is intertemporal *social* welfare rather than individual or household welfare. The DU approach leads to an SDR that is expressed by the so-called Ramsey rule, named after Frank Ramsey, an eminent mathematician and economist from the early twentieth century: the SDR equals *utility discounting* (expressed as the Greek letter rho, ρ), added to the *wealth effect*, which is a measure of aversion to inequality (expressed as the Greek letter eta, η) multiplied by *income growth* (g). Thus, in the form of an equation, $SDR = \rho + \eta g$. In this equation, ρ is known as the *pure rate of time preference* or *utility discount rate*, and it reflects the first reason for discounting the future: discounting future utilities. If $SDR = \rho + \eta g$, it's easy to see that the SDR increases as any one of its components—the utility discount rate, growth, or inequality aversion—increases.

That may all sound rather abstract and stylized, but precisely those principles appear in government guidance on CBA throughout the world, and in the IPCC's *Fifth Assessment Report*, Working Group III, they represent one of the central ways of thinking about how to evaluate climate change.[18] The US government's report establishing a value for climate change damages discusses at length how to evaluate intergenerational decisions. Here, table 1 shows how other governments around the world calibrate the Ramsey rule for use in their domestic CBAs, along with analogous approaches from important reports on climate change.

Table 1 and the discussions in the IPCC and US reports suggest that 3 percent isn't an unreasonable choice for an SDR. However, if we want to examine critiques of setting the SDR at 3 percent (recall that this is the rate at which the typical climate change project would not pass a CBA test)—particularly critiques that would place higher weight on future welfare—we must turn to the evidence underlying the parameters that go into an SDR.

Estimating the Social Discount Rate

Where do the numbers in table 1 come from? Let's consider each of the parameters in turn. The essence of the pure time preference, ρ, can be understood by first thinking about

Table 1. How Governments and Reports Calibrate the Ramsey Rule

Country/ Study	Pure time preference (ρ)	Inequality aversion (η)	Growth (g)	Social discount rate (SDR)	Source
United Kingdom	0.5% (1%)	1%	2%	3.5% (1%)	HM Treasury (2003)
France	0%	2%	2%	4% (2%)	Lebègue (2005)
Stern	0% (0.1%)*	1%	1.3	1.4%	Stern Review (2007)
IPCC	0%	1–2%	2%	2–4%	IPCC (2013)
Nordhaus	2–3%	1%	2%	4–5%	Nordhaus (2007)

Notes: The rates in parentheses for the United Kingdom and France are the rates of discount for time horizons longer than 300 years. In the French case, the reduction occurs at 30 years. In the UK case, there is a stepped decline from 3.5 to 1 percent over that period.

Sources: HM Treasury, *The Green Book: Appraisal and Evaluation in Central Government* (London: HM Treasury, 2003; revised 2011); Intergovernmental Panel on Climate Change, *Climate Change 2013: The Physical Science Basis* (Cambridge: Cambridge University Press, 2013); Daniel Lebègue, *Révision du Taux d'Actualisation des Investissements Publics* (Paris: Commissariat Générale du Plan, 2005); William D. Nordhaus, "A Review of the *Stern Review on the Economics of Climate Change*," *Journal of Economic Literature* 45 (2007): 686–702, doi: 10.1257/jel.45.3.686; Nicholas Stern, *The Economics of Climate Change: The Stern Review* (Cambridge: Cambridge University Press, 2007).

*The figure in parentheses reflects the likelihood that society won't exist because of some catastrophic event, which Stern added to the pure time preference of 0 percent. The Stern SDR is an average of many different analyses contained in the *Stern Review*.

impatience. For instance, research has demonstrated empirically that children aren't always particularly good at deferring gratification in relation to things like, say, marshmallows.[19] They prefer to get their utility now rather than at even some very short time in the future: they are very impatient for marshmallows. That's an example of impatience at the individual level: individuals making their own decisions for their own benefit.

To evaluate societal projects, we need a measure of impatience that's appropriate for society as a whole. It must reflect the fact that decisions have implications not just for today's society but also for future, as yet unborn, generations. Some people argue that observing how people behave is the correct way to get such information. Others, particularly in the context of climate change, argue that consulting ethical principles is more appropriate.

In the context of climate change and of long-run CBA in general, the pure rate of time preference is typically treated as a normative parameter, to be guided by ethical arguments. The utilitarian tradition argues for treating generations equally—that is, $\rho = 0$—on the ethical grounds that we should be impartial about when a person is born or when a society exists. That is, societies should be *anonymous*. The consequence of doing otherwise would be that generations in the distant future would be tyrannized, in the sense that a weight of zero would be placed on their utility and hence in the DU measure of intertemporal social welfare.

Just as there are ethical arguments for setting pure time preference at zero, there are ethical arguments for ρ greater than zero.[20] Nobel Prize–winning economist

Kenneth Arrow has discussed the trade-off between morality and self-regard as an ethical argument for treating current and future generations' utilities unequally.[21] His argument is that imposing equal treatment may tyrannize the present through onerous savings or investment requirements. In essence, when the utility discount rate is zero, increments to the utility of generations millions of years into the future have the same effect on social welfare as do increments today. Moreover, there are many of those generations! This would indicate that the current generation ought to invest in many more lower-return projects. However, the notion of self-regard proposes that individuals need not adhere to the morality of equal treatment if it comes at too great a cost to themselves. In particular, Arrow concludes that "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."[22]

Another argument for discounting future utilities is the possibility that as a result of some external catastrophe, future generations may not exist at all. In that case, a separate term would be added to the pure rate of time preference to reflect the hazard rate of catastrophe.[23] This argument has found some support within the utilitarian tradition, although recently, some detractors have said that "from the ethical standpoint it may . . . seem questionable to make such a bet on the existence of future generations."[24]

How have recent studies of climate change approached the issue of where to set the value of ρ? The *Stern Review*—a highly influential report on the economics of climate change undertaken for the UK government—took the view that barring a small probability of global societal collapse of 0.1 percent per year, each generation's wellbeing should be treated equally: thus, $\rho = 0$ percent, but to that should be added a hazard rate of 0.1 percent. This choice contrasts with the UK Treasury guidelines on cost–benefit analysis referred to in table 1. Based on a variety of empirical studies, these guidelines argue that the risk of catastrophe in the UK is 1 percent per year. On top of that is a pure time preference of 0.5 percent, leading to an overall discount rate for utility of 1.5 percent.

Nordhaus, the developer of the DICE model, took a very different approach. He made the additional assumption that markets would equate the SDR in the equation SDR $= \rho + \eta g$ to the market rate of interest. He then calibrated the parameters (particularly ρ) around the market interest rate, using empirical estimates of g and η.[25] That led to ρ of 1.5 to 3 percent. Linking the calculation to observed market rates is sometimes referred to as a *positive* or *descriptive* approach to identifying the correct SDR. And we've covered but a few of the ways to estimate the pure rate of time preference.[26]

Within the Ramsey framework of SDR = $\rho + \eta g$, the parameter η reflects aversion to income inequality. Here we're thinking about potential inequality across generations. However, there are other interpretations of the parameter in different contexts. For instance, it might also be assumed to govern inequality aversion between individuals at the same point in time or inequality aversion across different risky states of the world.[27] Consequently, people have used different methods of estimating parameter η: for example, progressivity of income tax schedules (known as intratemporal inequality aversion), ethical introspection (intratemporal or intertemporal inequality aversion), international transfers of aid (international inequality aversion), observed consumption behavior at the aggregate or individual level (intertemporal substitution), experiments involving risk (risk aversion), and so on.

In the UK case, evidence to guide estimates of η comes from a variety of those sources.[28] The most recent estimates for the UK from observed behavior tend to suggest a value of around 1.5 to 1.6, whatever the type of data used.[29] Several experts have suggested a value of 2 on the basis of ethical considerations and personal introspection.[30]

Taken together, different perspectives on the parameters of the Ramsey rule naturally lead to different recommendations for an SDR. With expected annual growth of 2 percent, the UK selections of ρ = 1.5 percent and η = 1 lead to an SDR of 3.5 percent. In France, expected growth of 2 percent together with ρ = 0 and η = 2 has led to an SDR of 4 percent (see table 1). The US analysis ultimately proposed SDR values of 2.5, 3, and 5 percent, with 3 percent being the central case around which the report undertakes sensitivity analysis. All of this suggests that an SDR of 3 percent, with its consequent weight of one-twentieth for net benefits a hundred years in the future, can be criticized simply by disagreeing with the interpretation of evidence and the ethical rationale. A number at the lower end of the given examples—say, 1.4 percent—would apply a weight of almost one-fourth to those net benefits a hundred years in the future, counting them almost five times as much against current costs as a 3 percent SDR would.

Discounted Utilitarianism Extended: Uncertainty about Growth

Even if we agree on an appropriate pure rate of time preference and a level of inequality aversion, the weight we place on future generations depends on the economic state in which we think our descendants will find themselves. In particular, the economic growth rate during the next hundred years and beyond is very uncertain, and the differences among the climate change scenarios in figure 1 illustrate that. How does the economic framework deal with uncertainty, and what are the implications for the way we account for future generations when we calculate CBA today?

Suppose our descendants will be faced with one of two possible states of the world at some point in the future—say, a hundred years from now. One is a "good" state, in which annual incomes are high, at $30,000, and the other is a "bad" state, in which annual incomes are only $10,000. A typical way to summarize today the welfare we expect in the future would be to simply take the average of the utilities associated with each state of the world. And a typical way of summarizing the welfare impact of an intervention that, say, raises income by one dollar would be to take the average change in welfare associated with each state of the world. This is the *expected utility*, which could then be added up over time to obtain our intertemporal welfare function, as in the DU approach.

How does this affect our evaluation versus simply using the average income of $20,000? It depends on another dimension of societal preferences, known as *prudence*, which refers to the idea that as the future becomes more uncertain regarding the best guess about income, the value of an additional dollar in the future increases. More uncertainty about income then leads to more savings; hence we call that effect prudence. Assuming prudence at the societal level, the effective SDR should be lower if we are uncertain about the state of the world that our descendants will inherit versus our best estimate of the average outcome. Moreover, the higher the level of uncertainty, the lower the SDR. Thus, the prudence effect would likely be higher over longer time horizons, where uncertainty about the effect of growth is greater. The idea that uncertainty is greater the further we look into the future actually justifies the use of a discount rate that is smaller for costs and benefits that occur in the future compared with today's costs and benefits: a *declining discount rate*.

The presence and degree of those prudence effects are, in general, determined by the same parameter that describes inequality (and risk) aversion, and for most reasonable values of that parameter, such aversion is often both present and large. Across countries, uncertainty about future growth would tend to justify a discount rate of less than 1 percent for long time horizons.[31]

There is no doubt that DU, declining discount rates, and related economic theories have been extremely influential in policy circles.[32] The guidelines of both the UK Treasury and the US Environmental Protection Agency have been heavily influenced by them.[33] The Norwegian government, too, in its advice on the time profile of discount rates that are to be applied to different time horizons (the term structure), refers to arguments about uncertainty over the rate of return to capital.[34]

Table 1 shows that several governments have made this theory and its close relations a central part of their CBA guidelines. In France, the SDR declines from 4 to 2 percent after 30 years. In the United Kingdom, the SDR declines steadily over 300 years from an initial 3.5 percent to 1 percent. The United States uses a lower discount rate of 2.5 percent for intergenerational projects and to evaluate the social cost of carbon, which is the current value of all future damages arising from an additional ton of carbon emissions today.

So, in theory, uncertainty about future income levels increases the weight we place on our descendants' wellbeing. In practice, such uncertainty has been shown to be important for long-term policy making.[35] If uncertainty were to justify using a rate of 2.5 percent rather than 3 percent over the next hundred years, as suggested by the US government's analysis, our weight for net benefits a hundred years from now would change from one-twentieth of today's value to almost one-twelfth. But is that the only omission that the standard DU approach makes in its parsimonious approach to intertemporal decision making? When we put monetary values on many of the benefits of mitigating climate change, we are making certain assumptions about the value future generations will place on avoiding human and natural impacts—a topic to which we now turn.

Discounted Utilitarianism Extended: Environmental Goods and Services

When we evaluate courses of action today that will affect the wellbeing of our descendants in the future, we estimate both the monetized costs and benefits in each year and the weights necessary to compare those costs and benefits across time. As we noted in the previous section, nonmarket benefits such as health and environmental amenities aren't easily measured, but they could be quite large. Moreover, if future generations, otherwise equivalent to ourselves, simultaneously face a denuded environment and poorer health, the value they put on improving those amenities could be even larger than the value we put on similar improvements today. The environment and health are important dimensions of wellbeing. Thus, to evaluate how changes in those amenities are valued in the future, it's important to understand how they evolve over time alongside income and consumption.

Consider the following extension to the previous example, which looked at consumption growth and its effects on how we valued an additional dollar of income. Now consider two generations—the present and the future—whose wellbeing now depends on consumption *and* a measure of environmental services. Both have the same income levels, say, $20,000 per annum, and both consume identically. They differ only in the environmental services they each enjoy. Suppose that environmental services decline over time so that our descendants have 50 million hectares of forested land compared with our 200 million hectares (one hectare is about two and a half acres). How many dollars of consumption would each generation give up for an additional hectare of land?

If, like consumption, the added welfare from additional units of environmental services declines as the amount available rises, then due to land's increasing scarcity, our descendants would probably be willing to give up more dollars of consumption for a hectare of land than would the current generation. So increments of environmental goods are worth more to our descendants than they are to us, and we would place different values on changes to the environment depending on when they happen in time. If that sounds a lot like the wealth effect that we discussed earlier, then it should—only this is an *environmental* wealth effect, where increasing scarcity has the opposite effect of raising, rather than lowering, the value of changes that occur in the future, all things equal.

But precisely how does this affect the valuation of our descendants' wellbeing? Empirical estimates suggest that the price of environmental goods could be rising at an annual rate 1 percent faster than consumption goods, an indication of their relative scarcity.[36] Evaluating our descendants' wellbeing in this way increases the prescribed urgency of climate mitigation policies when compared with other analyses.[37] For example, if along with our 3 percent discount rate we undervalued future benefits at 1 percent per year, it would be appropriate to use a 2 percent discount rate, thus weighting a mismeasured future net benefit in a hundred years at almost one-seventh of today's value rather than one-twentieth.

Ultimately, to properly evaluate how courses of action would affect our descendants, we must be careful to use a set of accounting prices in our CBA that reflect the relative scarcity of environmental or health goods in the future, rather than assuming that today's amenity values will remain the same. Differences in relative prices can be equivalently reflected in adjustments to the discount rate, downward for increasing scarcity of environmental resources. We need further work to understand whether current approaches have gotten that right.

Alternatives to Discounted Utilitarianism and Uncertainty

So far, we've focused on reasons that net benefits to future generations might be undervalued while staying within the standard welfare framework. But is this the only way to look at the problem? Does DU satisfy all the tenets of fairness that we might want to satisfy when taking a position on our descendants' wellbeing? That question is particularly relevant given uncertainty about the future in terms of both how the economy will evolve and the potential for calamitous climate change. Absent uncertainty, welfare analysis—whether or not it considers fairness—tends to be dominated by the general expectation that future generations will be much wealthier than our own and that climate change impacts will only put a dent in the degree to which they are wealthier but will not alter the general trend. In that landscape, a sharper focus on fairness would tend to disfavor those wealthy future generations.

With uncertainty, we confront the real possibility—whether small or not—that the future could be worse for future generations than for ourselves. In that landscape, a sharper focus on fairness could *favor* future generations. What might a sharper focus on fairness look like? Naturally, there are alternatives to DU both within and outside economics. In recent years, particularly since the *Stern Review* appeared, interest in different conceptions of intertemporal welfare has been growing. Following are some examples to illustrate the point.

One approach is to simply increase aversion to inequality—specifically, aversion to the inequality that might occur for a future impoverished generation. In the DU framework, we often consider individual preferences to be appropriate sources of information about societal preferences, as we saw in our earlier discussion of how the parameters might be estimated. By simply asserting more inequality aversion or, somewhat equivalently, by adding aversion to unequal utility, we would raise the weight (relative to our own) placed on impacts for future generations that are worse off than our own. Prioritarianism, for instance, is an alternative to utilitarianism wherein, for reasons of fairness, the utility of generations that have the lowest utility levels receives more weight.

Of course, increased aversion to inequality also means lowering the weight (relative to our own) that we place on impacts for future generations that are better off than ours. Therefore, another concept would be a more nuanced form of increased aversion to inequality. Specifically, we could be averse to leaving a future generation worse off but not averse to future generations being better off. That concept pertains to the notion of sustainability. The sustainable discounted utilitarian approach is one example of how sustainability can be included in the analysis. Models using that approach have shown that taking sustainability into account could raise the level of willingness to pay for climate mitigation severalfold, reducing the effective SDR.[38]

A third possibility comes from recognizing that within the DU framework, aversion to inequality applies equally to inequality across time and generations as well as to inequality across risky outcomes (risk aversion). Recently, researchers have explored separating those two concepts when it comes to climate change.[39] Doing so lets us consider societal preferences that are more averse to climate risks—thus raising the value that future generations place on avoiding those risks—while maintaining the same relative weights between current and future generations based on average economic growth and/or the passage of time.

Finally, quite distinct from the question of uncertainty and aversion to risk and inequality, some researchers have begun to explore the question of population ethics related to climate change.[40] In most integrated assessment models, DICE included, mortality impacts are reflected by costing out lives through a method called *value of statistical life*. As the term suggests, value of statistical life is a statistical estimate, using observed or hypothetical behavior, of an individual's willingness to pay to reduce the risk of death.[41] Does that make sense, particularly when applied to the risk of large, catastrophic population impacts? To what extent are more people better than fewer if the larger population is worse off? Is there a critical, minimum level of utility below which life is not worth living?[42] Those and other questions are the topics of a new and evolving area of research, which presumably could lead to placing more weight on the consequences for future generations.

Conclusions

As we said at the outset, choices about climate change mitigation involve a tricky balance between the interests of current and future generations. Current generations largely bear the cost of mitigation; future generations largely reap the benefits—though, at the same time, they face similar trade-offs with their own future generations. To the extent that we're interested in how climate change affects children, it's hard to get away from the fact that today's children will grow up to be a future generation, as will their children. For them, how we make intergenerational trade-offs is likely to mean as much as—and perhaps more than—how we modify the estimated costs and benefits for adults so that those costs and benefits are instead appropriate for children at the same moment in time.

For that reason, we've explored what economic analysis can tell us about the balance of those costs and benefits and why our future children might criticize that analysis. Using a well-known model of climate change mitigation costs and benefits, we estimated that avoiding a 2°C temperature increase would not (quite) pass a cost–benefit analysis. We based our estimate on a discount rate of 3 percent, the value recently suggested by the US government. Such a discount rate implies that monetized benefits a hundred years in the future receive a weight of about one-twentieth of the weight given to monetized costs today.

Why might that weight be wrong? We've explored three main reasons. First, people disagree considerably about the correct discount rate. Other plausible interpretations of society's preferences or observed data would increase that weight from one-twentieth to one-quarter—a factor of five. Even using the standard parameters but acknowledging that future economic growth is uncertain could change the weight to almost one-twelfth.

Second, we may have failed to correctly value future climate change impacts, particularly those involving the loss of environmental amenities that have no close monetary substitutes. One calculation suggests that accounting for those impacts might mean adjusting a future weight from one-twentieth to one-seventh.

Finally, we also examined how uncertainty and alternatives to the standard welfare approach might affect future valuation. Here the work is more recent and more speculative. However, properly valuing catastrophic risks, and particularly the risk of major population changes, could alter the way we value impacts on future generations.

Ultimately, our goal has not been to provide a different or better answer to the question of how we should value future climate change impacts on our children. Instead, we have tried to explain how current economic analysis treats our children and our children's children in terms of intergenerational welfare. We've also tried to explain why current economic analysis might be wrong and, when possible, by how much, focusing on why an error might undervalue the future.

None of this should be interpreted as a failure of current analysis: policy making at any moment in time requires the best information and judgment available. We believe current efforts to balance climate change costs and benefits are valuable, though they may be only part of the answer. Ultimately, many of the choices come down to ethical questions, and many of the decisions come down to societal preferences—all of which will be difficult to extract from data or theory.

ENDNOTES

1 . See William D. Nordhaus, "Rolling the 'DICE': An Optimal Transition Path for Controlling Greenhouse Gases," *Resource and Energy Economics* 15 (1993): 27–50, doi: 10.1016/0928-7655(93)90017-O; William D. Nordhaus, *Question of Balance: Weighing the Options on Global Warming Policies* (New Haven, CT: Yale University Press, 2008); and William D. Nordhaus, "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches," *Journal of the Association of Environmental and Resource Economists* 1 (2014): 273–312, doi: 10.1086/676035. We acknowledge Louise Kessler's research assistance in running the DICE model.

2. Simon Dietz, Christian Gollier, and Louise Kessler, "The Climate Beta," working papers no. 215 [Centre for Climate Change Economics and Policy] and no. 190 [Grantham Research Institute on Climate Change and the Environment] (London School of Economics and Political Science, London, 2015).

3. Interagency Working Group on Social Cost of Carbon, US Government, "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866" (May 2013), https://www.whitehouse.gov/sites/default/files/omb/inforeg/social_cost_of_carbon_for_ria_2013_update.pdf.

4. Michael E. Porter and Claas van der Linde, "Toward a New Conception of the Environment–Competitiveness Relationship," *Journal of Economic Perspectives* 9, no. 4 (1995): 97–118, doi: 10.1257/jep.9.4.97; Richard D. Morgenstern, William A. Pizer, and Jhih-Shyang Shih, "The Cost of Environmental Protection," *Review of Economics and Statistics* 83 (2001): 732–38, doi: 10.1162/003465301753237812.

5. Tomas Nauclér and Per-Anders Enkvist, *Pathways to a Low-Carbon Economy: Version 2 of the Global Greenhouse Gas Abatement Cost Curve* (New York: McKinsey & Company, 2009).

6. Simon Dietz and Cameron J. Hepburn, "Benefit–Cost Analysis of Non-Marginal Climate and Energy Projects," *Energy Economics* 40 (2013): 61–71, doi: 10.1016/j.eneco.2013.05.023; Massimo Tavoni and Richard S. J. Tol, "Counting Only the Hits? The Risk of Underestimating the Costs of Stringent Climate Policy: A Letter," *Climatic Change* 100 (2010): 769–78, doi: 10.1007/s10584-010-9867-9.

7. Valentina Bosetti, Carlo Carraro, and Massimo Tavoni, "Timing of Mitigation and Technology Availability in Achieving a Low-Carbon World," *Environmental and Resource Economics* 51 (2012): 353–69, doi: 10.1007/s10640-011-9502-x.

8. Elmar Kriegler et al., "The Role of Technology for Achieving Climate Policy Objectives: Overview of the EMF 27 Study on Global Technology and Climate Policy Strategies," *Climatic Change* 123 (2014): 353–67, doi: 10.1007/s10584-013-0953-7.

9. Global CCS Institute, "Large Scale CCS Projects" (2015), http://www.globalccsinstitute.com/projects/large-scale-ccs-projects.

10. Christoph Böhringer, Edward J. Balistreri, and Thomas F. Rutherford, "The Role of Border Carbon Adjustment in Unilateral Climate Policy: Overview of an Energy Modeling Forum Study (EMF 29)," *Energy Economics* 34 (2012), S97–S110, doi: 10.1016/j.eneco.2012.10.003.

11. Gregory F. Nemet, Tracey Holloway, and Paul Meier, "Implications of Incorporating Air-Quality Co-Benefits into Climate Change Policymaking," *Environmental Research Letters* 5 (2010): 014007, doi: 10.1088/1748-9326/5/1/014007.

12. Sergey Paltsev et al., "How (and Why) Do Climate Policy Costs Differ among Countries?" in *Human-Induced Climate Change,* ed. Michael E. Schlesinger et al. (Cambridge: Cambridge University Press, 2007), 282–93.

13. William D. Nordhaus and Joseph Boyer, *Warming the World: Economic Models of Global Warming* (Cambridge, MA: MIT Press, 2000).

14. Richard S. J. Tol, "Correction and Update: The Economic Effects of Climate Change," *Journal of Economic Perspectives* 28 (2000): 221–26, doi: 10.1257/jep.28.2.221.
15. Robert Mendelsohn and Michael E. Schlesinger, "Climate-Response Functions," *Ambio* 28 (1999): 362–66.
16. Martin Weitzman, "GHG Targets as Insurance against Catastrophic Climate Damages," *Journal of Public Economic Theory* 14 (2012): 221–44, doi: 10.1111/j.1467-9779.2011.01539.x.
17. Paul Watkiss, "Aggregate Economic Measures of Climate Change Damages: Explaining the Differences and Implications," *Wiley Interdisciplinary Reviews: Climate Change* 2 (2011): 356–72, doi: 10.1002/wcc.111.
18. IPCC, *Climate Change 2013: The Physical Science Basis* (Cambridge: Cambridge University Press, 2013).
19. Walter Mischel, Yuichi Shoda, and Monica L. Rodriguez, "Delay of Gratification in Children," *Science* 244 (1989): 933–38, doi: 10.1126/science.2658056.
20. Wilfred Beckerman and Cameron Hepburn, "Ethics of the Discount Rate in the *Stern Review on the Economics of Climate Change*," *World Economics* 8 (2007): 187–210.
21. Kenneth J. Arrow, "Discounting, Morality, and Gaming," in *Discounting and Intergenerational Equity*, ed. Paul R. Portney and John P. Weyant (Washington, DC: Resources for the Future, 1999), 13–21.
22. Ibid.
23. The mechanics of this process are discussed in Partha S. Dasgupta and Geoffrey M. Heal, eds., *Economic Theory and Exhaustible Resources* (Oxford: Cambridge University Press, 1979), 260–65.
24. Wolfgang Buchholz and Michael Schymura, "Intertemporal Evaluation Criteria for Climate Change Policy: The Basic Ethical Issues," Discussion Paper no. 11-031 (Centre for European Economic Research [ZEW], Mannheim, Germany, 2011), 16.
25. William D. Nordhaus, "A Review of the *Stern Review on the Economics of Climate Change*," *Journal of Economic Literature* 45 (2007): 686–702, doi: 10.1257/jel.45.3.686; William Nordhaus, *A Question of Balance: Weighing the Options on Global Warming Policies* (New Haven, CT: Yale University Press, 2008).
26. A recent survey of 200 experts in the context of long-term CBA shows the extent of disagreement on the utility discount rate. The minimum value was zero, which was also the modal value. The maximum value was 8 percent. For details, see Moritz Drupp et al., "Discounting Disentangled: An Expert Survey on the Determinants of the Long-Term Social Discount Rate," Working Paper 195 (Centre for Climate Change Economics and Policy, London School of Economics and Political Science, London, 2015).
27. For more on these concepts and how to measure inequality aversion, see Nicholas Stern, "The Marginal Valuation of Income," in *Studies in Modern Economic Analysis: The Proceedings of the Association of University Teachers of Economics, Edinburgh, 1976*, ed. Michael J. Artis and A. Robert Nobay (Oxford, UK: Basil Blackwell, 1977), 209–54; and Dasgupta, "Discounting." Empirical evidence suggests that individuals, even when considering intergenerational contexts, don't equate these concepts. See Håkon Sælen et al., "Siblings, Not Triplets: Social Preferences for Risk, Inequality and Time in Discounting Climate Change," *Economics: The Open-Access, Open-Assessment E-Journal*, 3 (2009–26): 1–28, doi: 10.5018/economics-ejournal.ja.2009-26.
28. See Stern, "Welfare Weights," and OXERA Consulting, "A Social Time Preference Rate for Use in Long-Term Discounting" (Oxford, UK, 2002).
29. Ben Groom and David Maddison, "Non-Identical Quadruplets: Four New Estimates of the Elasticity of Marginal Utility for the UK," Working Paper 121 (Centre for Climate Change Economics and Policy, London School of Economics and Political Science, London, 2013), looks at estimates of the elasticity of intertemporal substitution via analysis of the consumption function, inequality aversion via progressive tax rates, and Ragnar Frisch's wants-independence approach.

30. See Drupp et al., "Discounting Disentangled."

31. This calculation comes from Christian Gollier, "On the Underestimation of the Precautionary Effect in Discounting," Working Paper 3536 (CESifo, Munich, 2011), on the basis of a cross-country analysis of growth volatility.

32. For example, Martin Weitzman, "Gamma Discounting," *American Economic Review* 91 (2001): 260–71, and Martin Weitzman, "Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate," *Journal of Environmental Economics and Management* 36 (1998): 201–8.

33. HM Treasury, *The Green Book: Appraisal and Evaluation in Central Government* (London: HM Treasury, 2003; revised 2011), and National Center for Environmental Economics, Office of Policy, US Environmental Protection Agency, *Guidelines for Preparing Economic Analyses* (Washington, DC: US Environmental Protection Agency, 2010).

34. There are issues concerning implementation and estimation, but these arguments have gained a lot of traction in policy. See Maureen L. Cropper et al., "Declining Discount Rates," *American Economic Review* 104 (2014): 538–43, doi: 10.1257/aer.104.5.538.

35. This is not an exhaustive list of justifications for a declining term structure of social discount rates. Indeed, many scholars point to the use of a declining term structure of discount rates for risk-free projects. For excellent surveys of the theory, empirics, and practice, see Kenneth J. Arrow et al., "Should Governments Use a Declining Discount Rate in Project Analysis?" *Review of Environmental Economics and Policy* 8 (2014): 145–63, doi: 10.1093/reep/reu008, and Christian Gollier, *Pricing the Future: The Economics of Discounting in an Uncertain World* (Princeton, NJ: Princeton University Press, 2012).

36. See Stefan Baumgärtner et al., "Ramsey Discounting of Ecosystem Services," *Environmental and Resource Economics* 61 (2015): 273–96, doi: 10.1007/s10640-014-9792-x.

37. Important references here include Thomas Sterner and U. Martin Persson, "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate," *Review of Environmental Economics and Policy* 2 (2008): 61–76, doi: 10.1093/reep/rem024; Michael Hoel and Thomas Sterner, "Discounting and Relative Prices," *Climatic Change* 84 (2007): 265–80, doi: 10.1007/s10584-007-9255-2; and Hans-Peter Weikard and Xueqin Zhu, "Discounting and Environmental Quality: When Should Dual Rates Be Used?" *Economic Modelling* 22 (2005): 868–78, doi: 10.1016/j.econmod.2005.06.004.

38. Higher willingness to pay is shown in Simon Dietz and Geir B. Asheim, "Climate Policy under Sustainable Discounted Utilitarianism," *Journal of Environmental Economics and Management* 63 (2012): 321–35, doi: 10.1016/j.jeem.2012.01.003.

39. For example, Ravi Bansal, Dana Kiku, and Marcelo Ochoa, "Climate Change and Growth Risks (paper presented at National Bureau of Economic Research 2015 Summer Institute, Cambridge, MA, July 20–21, 2015), http://www.nber.org/confer//2015/SI2015/EEE/Bansal_Kiku_Ochoa.pdf, shows that such a model implies a significant cost of climate damages.

40. Antony Millner, "On Welfare Frameworks and Catastrophic Climate Risks," *Journal of Environmental Economics and Management* 65 (2013): 310–25.

41. For a review of value of statistical life in the context of climate change, see Maureen L. Cropper, "Estimating the Economic Value of Health Impacts of Climate Change" (paper presented at Research on Climate Change Impacts and Associated Economic Damages workshop, Washington, DC, January 27–28, 2011), http://yosemite.epa.gov/ee/epa/eerm.nsf/vwAN/EE-0566-110.pdf/$file/EE-0566-110.pdf.

42. See Marc Fleurbaey and Stéphane Zuber, "Discounting, beyond Utilitarianism," Economics Discussion Papers 40 (Kiel Institute for the World Economy, Kiel, Germany, 2014).