

Polymaking for posterity

Lawrence Summers · Richard Zeckhauser

Published online: 30 September 2008
© Springer Science + Business Media, LLC 2008

Abstract Polymaking for posterity involves current decisions with distant consequences. Contrary to conventional prescriptions, we conclude that the greater wealth of future generations may strengthen the case for preserving environmental amenities; lower discount rates should be applied to the far future, and special effort should be made to avoid actions that impose costs on future generations. Posterity brings great uncertainties. Even massive losses, such as human extinction, however, do not merit infinite negative utility. Given learning, greater uncertainties about damages could increase or decrease the optimal level of current mitigation activities. Policies for posterity should anticipate effects on: alternative investments, both public and private; the actions of other nations; and the behaviors of future generations. Such effects may surprise. This analysis blends traditional public finance and behavioral economics with a number of hypothetical choice problems.

Keywords Discounting · Posterity · Altruism · Comet problem · Trolley problem · Climate change · Global warming · Uncertainty · Learning · Reaction function

JEL Classification D90 · D64 · Q54 · D81

Comparisons of present and future benefits and costs are ubiquitous in the formulation of public policy. How much research should governments support, what lifetimes should they seek for their buildings and monuments, what levels of transfers should they make to children and to the aged, how should they think about using debt to finance their activities, and how should they value the future environmental consequences of actions taken today? To answer these questions requires comparisons of present and future benefits and costs.

L. Summers · R. Zeckhauser (✉)
Harvard University, Cambridge, MA, USA
e-mail: richard_zeckhauser@harvard.edu

A vast literature has sought to provide ethical and economic arguments as to how such comparisons can best be made and can at least in principle take account of issues such as uncertainty, the difficulties the government has in mobilizing resources, inefficiencies from taxation, the alternative private use of resources, and so forth. A focus is often placed on the appropriate discount rate for the government to use in evaluating various kinds of projects. Frederick et al. (2002) surveys this literature, and Portney and Weyant (1999) distills a number of thoughtful analyses regarding valuation of the future. Two widely held conclusions are that market rates of discount are not the appropriate indicator when considering projects undertaken on behalf of the society as a whole and that market rates are well above society's true discount rate. The implication is that social investments should weight the future far higher than do private projects.

An array of critics, notably including philosophers and scientists, point out that the discounting approach routinely ignores a host of ethical issues. For example, future generations have legitimate interests, and those interests deserve more weight than the altruistic concerns of those presently alive would give them. A second strand of ethical argument emerges from a rights literature. This argument, at least implicitly, underscores much of the support for strong action against global warming, or more generally climate change. Human beings, most notably those currently alive, have taken actions that have despoiled and continue to despoil the environment. The future has as much right as we did to inherit a hospitable Earth. In some sense this is a generalization of the admonition one sees when walking into a nature area: "Leave this natural habitat as you found it." This argument would suggest that if society's choice were between spending an incremental \$1 billion dollars curbing environmental damage or an equivalent amount on medical research, even if all future generations would be better off with the latter, the resources should go to curbing the damage.

This paper addresses policymaking for *posterity*. We employ this term to invoke intervals of several decades or centuries, not mere years. For a variety of reasons, starting with the observation that over the very long run those alive today will not be alive to bear the very long-run consequences of their actions, we believe that very long-run policymaking needs to be conceptualized differently than policymaking over shorter periods. Equally important, we can never know what the future holds, but for the far future the uncertainties are massive.

The posterity problem is of considerable practical importance, since many contemporary policy choices will have costs and benefits of policy that play out over very long periods of time. Very long horizons are involved in decisions regarding expensive extremely fundamental research, the gradual reform of entitlement programs for the elderly, the storage of nuclear wastes, the avoidance of nuclear proliferation, and the protection of the planet from asteroids and comets. But distant horizons have received the greatest policy attention in relation to climate change. As the celebrated Stern (2007) Report illustrates, much of the argument regarding global climate change turns on the consequences a century or more hence from actions taken or not taken today.

Nordhaus's (2007) commentary on Stern demonstrates that conclusions regarding long-run discount rates have an enormous impact on policy judgments regarding

climate change. He finds that Stern's very strong conclusions regarding the need for major immediate actions to address this area are not driven by new views regarding either its consequences or the cost of preventing it but by his assumptions regarding discount rates. This should not be surprising. At even a relatively modest 3% discount rate—a dollar of benefits a century from now is worth less than 6 cents today and a dollar two centuries from now is worth less than 0.4 cents. At the discount rate of 7% mandated for use in certain US government contexts by the OMB, the distant future becomes nearly irrelevant, as \$100 a century from now is valued less than 10 cents today.

Our goal is not to reach an ultimate conclusion regarding “the social discount rate.” Rather it is to contribute to better policymaking for the very long run by identifying a range of considerations that should enter into society's judgments about taking actions that will have important consequences for our relatively distant descendants. We use global climate change as our motivating example throughout. Our analysis proceeds in four sections.

Section 1 examines very long-run policymaking in the context of certainty. It argues that even if there is in general a reasonably compelling case for the use of a significantly positive discount rate in evaluating most public policy questions, there are a number of considerations that suggest giving significant weight to the distant future in considering issues like global warming. These derive from reflection on a number of aspects of the preferences that people reveal in experiments and in their everyday behaviors, including their undifferentiated views of posterity, their disproportionate aversion to suffering or causing losses, and the tendency for the value of non-pecuniary goods to rise disproportionately with increases in income. Our findings in this section generally suggest taking more action on behalf of the future than would the standard spare discounting approach.

Section 2 takes up issues relating when uncertainty is a central concern, as it surely is with most issues that stretch across long periods, and for sure with climate change. Many commentators—notably Weitzman (2007) and advocates of the so called precautionary principle—view the tremendous uncertainty regarding global climate change as an important argument for accelerating action, or what leads to the same actions, using a very low discount rate in measuring the benefits of climate change mitigation policies. We conclude, after careful consideration of how to model preferences when catastrophes threaten, and from recognizing that uncertainty introduces the possibility of future learning, that the pure uncertainty case for giving greater weight to future risks is less compelling than many have suggested.

Section 3 takes up what we label “reaction function issues.” Any choice that society makes today has consequences for other choices that private actors and subsequent societies can and will make in the future. That effect is a crucial consideration in policymaking for the very long run. While such reaction function issues are always present in policymaking, we identify a number of reasons why they bear in particular on how projects with very long consequences should be evaluated. Other reaction-related issues addressed include generational turnover in the identity and possible preferences of policymakers, the global aspect of very long-run issues, and questions of how priorities are balanced across areas.

Section 4 provides a summary, and suggests some directions for future work.

1 Policy for posterity given certainty

Economists' standard approach to comparing future costs and benefits is reflected in what might be labeled the fundamental discounting equation:

$$\rho = \delta + g * \eta,$$

where ρ is the discount rate to be applied in valuing future per-capita consumption dollars, δ is taken as a measure of pure time preference, g is the rate of growth of per capita income, and η is minus the elasticity of marginal utility with respect to consumption. In essence the equation says that a dollar in the future is worth less than a dollar for two reasons: First, the future increments to consumption are discounted because of pure impatience. The δ term would also reflect the fact that we give less weight to those alive in the future than we give to ourselves. The second term shows that increments to future consumption are reduced in value because the future will be richer, implying that an increment to consumption will be worth less in terms of marginal utility.

There has been much discussion of the appropriate assumptions to make regarding each of these parameters, and of consequences of uncertainty regarding them. The discussions quickly become philosophical. For example—is δ to be thought of as reflecting the value that those alive today place on the welfare of future generations or to represent the values that an imagined impartial ethical observer, with no particular generational connection, places on different generations? Weitzman (2007) suggests as an easy-to-remember triad of values $\delta = .02$, $g = .02$ and $\eta = 2$ which together imply a discount rate of 6% a year, a sufficiently high value to render anything that happens a century from now almost irrelevant as a dollar then is worth less than \$0.03 today [$0.03 > 0.0029472262 = 1/(1.06)^{100}$] (see also Cowen 2008). Others, such as Stern (2007), argue for alternative and much lower parameter values.

We do not believe that an argument about these alternative parameter values will resolve our policymaking concerns. Hence, we focus on several issues that we regard as important in making very long-run policy judgments that do not fit naturally into the standard approach.

Global population issues The fundamental discounting equation simply ignores questions relating to the size of the population. We are inclined to believe that the weight attached to the utility of a given generation should be related to its size, and probably proportional. In this case it is necessary to subtract the rate of population growth in discounting future levels of consumption. At the global level or for the USA such a correction would not be inconsequential, as population may well grow at close to 1% a year for the next half century.

Population issues raise further critical questions. When making policy for the very long run, how should a country think about those now living abroad who will immigrate and become citizens in the future? How about those whose descendants will become citizens? How about those nationals whose descendants will marry and have children with noncitizens?

In an increasingly open world it seems reasonable to suppose that the longer the horizon the more cosmopolitan the perspective needs to be, and therefore the size of

the population one should care about should increase. To see this point think about a Boston city councilor who cares only about the welfare of Bostonians. She would give little weight to some benefit Boston might provide non-Boston residents of Massachusetts next year. On the other hand if considering benefits that would only be realized a half century from now, the distinction between Boston and the state seems much less consequential, making the case for a broader perspective and other things equal a lower discount rate.

When this argument is extrapolated to the world at large, the issue is not merely that people will move around and intermarry. Some populations will grow much faster than others. How will Europeans or Japanese, or indeed most currently prosperous ethnic groups, weight the future if they conclude (as at least present evidence would suggest) that their percentage of the far-future world population is likely to be much smaller than it is today? We provide no ethical answer, but we would observe that this factor would lead them to weight the future lower.

Consequences of income growth The argument is sometimes made that the future will be far richer than are we, and that therefore providing for them by providing a less tarnished planet makes little sense; it is redistribution in the wrong direction. In essence, this is giving attention to the second term in the discounting equation above.

Of course the argument is not that nothing should be done to provide for the future. If this were done it is unlikely that growth would be positive. Rather the argument is that the distant future is so great a beneficiary of the spillover from investments made for the near term future that no special efforts on its behalf are appropriate. When the Greatest Generation made enormous sacrifices to fight World War II, they were preserving liberties in the twenty-first century. But they had sufficient motivation to preserve liberty in the 1940s and 1950s. When Thomas Edison pursued his extraordinary array of inventions, he effectively moved all successive worlds years ahead in their technological development. Thus, the Internet might have been delayed for a decade were it not for him. But Edison was substantially pursuing his personal interests, and we are just a lucky beneficiary. And so it is with all commercial R&D work today. The future benefits because much of the information generated is not appropriable.

The standard discounting approach provides the right way to think about the trade offs involved in helping the distant future if what is being traded off is current consumption that is perfectly substitutable with consumption in the distant future. In fact, in most policy contexts the benefits to the future are provided *in kind*, for example in the form of a protected rather than a despoiled environment.

In the environmental case, the benefit to the future of costly actions that we take today needs to be measured in terms of the future's willingness to pay for the amenity, a factor that is usually overlooked. And this will rise with their income, thereby counterbalancing the diminishing marginal utility of consumption as a valuation of consumption. In general, the effect of income growth on the value of providing a future amenity will be ambiguous. Consider the particular special case where amenity value is separable, so that a generation's utility can be written as $U(c) + a$, where c is consumption and a is the amenity. Then policymakers can consider amenity values by using current willingness-to-pay measures and then discounting

only at the pure rate of time preference whether coming from impatience or generational selfishness.¹

How important is this amenity elasticity issue? Potentially quite important. We would claim on the basis of casual empiricism that not only are societies prepared to pay more for environmental amenities or good health, for example, as they get richer but they are also more willing to sacrifice growth itself. More generally issues of health and the environment certainly loom much larger in rich than in poor countries.

As we roll into the future, if the past provides any prologue, production goods will be far cheaper. Looking at what might be called terms of trade, amenity goods will be far more valuable, and doubly so if—as we might expect—the income elasticity of demand for such goods is high. We could see people in the far future paying many dozens of times as much as people would today in terms of dishwashers or televisions sacrificed for an authentic wilderness experience, or a magnificent 70° spring day in New York.

Characterizing altruism When considering time preference in the context of altruistic preferences, it is hardly obvious that the standard discounting approach applies. Following philosophers, we will consider a hypothetical ethics problem to help get our thinking straight.

There are two giant comets that will be coming near to the Earth in the next decade, and will return again in the far future. Astronomers calculate that comet A has one chance in 100 of hitting the Earth on its next swing by, which will come in 100 years. Comet B has one chance in 10 of hitting the Earth on its next swing by, which will come in 400 years. If either comet is on course, there will be nothing that can be done at its next coming. Either comet will wipe out all sentient life if it hits. In preparation for killer comets, the government has prepared a single missile-weapon combination that can knock a comet significantly off course. Should it be fired at comet A or comet B?

Which comet would you target? Our intuition, and we suspect that of most of our readers, is that we would target Comet B. Yet a policymaker applying any realistic discount rate in the standard framework would choose Comet A using any discount rate even close to or above 1%.² But we would argue that it is not clear that we care about people 300 years from now more than modestly less than those alive 100 years from now. True, destroying A would protect an additional increment of people, namely those living from 100 to 300 years from now, but that hardly would take care of a 10 to 1 disparity in risk. We regard this as at least a tough conundrum. Moreover, unlike many discounting problems, it would not matter much to us how rich these various groups would be. Talk of great future affluence is a dodge often used when not providing for the future.

The comet problem convinces us that when we think about people in the far future, we should try to think hard about something like a weighting factor on their

¹ If the income elasticity of demand for the amenity exceeds the absolute value of the elasticity of marginal utility with respect to consumption, this approach would be an underestimate of how to value an amenity to a future generation, and vice versa. This formulation is simplifying and assuming a generation lives but one period.

² With a 1% discount rate, delaying a loss by 300 years cuts its value to 7.7% of its original value.

welfare. Traditional discounting analysis does not help much. And that weighting factor, at least for many of us, will be far greater than any discounting analysis would provide. Survey results from Cropper et al. (1994) indicate that individuals do attach weights to the welfares of far future generations well above what a discounting approach with any plausible interest rate could accommodate.

However, such generous weighting leads to a new conundrum when we recognize that the future potentially involves an infinite number of generations. If beyond some point weights to future generations are not going to decline, then a loss that continues forever, whenever it starts and however small, will count infinitely against benefits today. As noted above, we are willing to give far higher weights to far future generations than would a typical discounting analysis, but we think it essential to have in place some system where a finite loss to each generation over an infinite future counts far less than infinity.³

Such an approach might involve discounting, or thinking about a representative future generation. Enthusiasm for giving nontrivial weighting to a specific far future generation, it seems, may represent some variant of the embeddedness phenomenon well known in the contingent valuation literature: “Different but similar samples of respondents are asked about their willingness to pay for prevention of environmental damage scenarios that are identical except for their scale: different numbers of seabirds saved, different numbers of forest tracts preserved from logging, etc. It is reported that average willingness to pay is often substantial for the smallest scenario presented but is then substantially independent of the size of the damage averted, rising slightly if at all for large changes in size” (Arrow et al. 1993, p. 26). In our context, asking about valuing a cost to a single generation at a far future date may lead to a valuation not much less than asking about that same cost to all generations starting from that far future date.

Reference points, loss aversion, and errors of commission Many observers will argue that climate change is special, that we have more of an obligation to avoid destroying something that has been bequeathed to us by nature than to provide the future with say enhanced intellectual capital (from R&D, say). There are three strong elements to this argument: reference points, loss aversion, and errors of commission. Prospect theory (Kahneman and Tversky 1979) tells us that utility does not attach to the state space, rather to changes from the state space. It tells us as well that individuals treat losses from some reference point as being much more consequential than gains.

Every time a climate-change agreement is formulated, there is a profound debate on where should we start, what the reference point should be. Should it be current emissions, 2006 emissions (the baseline for the July 2008 Harvard University plan to curtail greenhouse gases by 30% by 2016), or some other date. When the G-8 countries agreed to a 50% cut by 2050, Japan wanted to start from 2008 levels; the Europeans favored 1990. The agreement specifies no date.

³ To be sure, merely assuming that there is a small probability that the world will end for some reason over which we have no control, for example a devastating war, can lead to effective declining weights as generations roll forward, but we are using this device to rescue the analysis.

Most discussions set a reference point for greenhouse gases, but if significant economic sacrifices are at stake, there are likely to be reference points for those losses as well. Prime ministers and presidents cannot hope to sacrifice GNP for climate control without hearing strong cries of protest from those experiencing loss aversion on their incomes. In growth economies, matters will be easier, because the cuts will come against what would have been their incomes, a moving and hence more fuzzy reference point. The “Save More Tomorrow” plan of Thaler and Benartzi (2001) increases employee savings by having them commit funds out of their future salary increases, presumably something that is less noticed, that serves less as a reference point. Our strong suspicion is that expenditures for climate change will be far easier to make in economies where per capita income is growing.

Environmental problems often involve an additional element: One’s purposeful action determines the change. That is where the contrast between errors of omission and commission play such a strong role. Consider the everyday problem of litter in a park. We all would prefer to encounter three rather than a dozen pieces of litter. The omission-commission distinction pushes further. Most citizens would never consider picking up some of those dozen littered pieces. But they would also never litter, whatever the current level. Thus, if a citizen inadvertently dropped a piece of paper she would pick it up, to avoid an act of commission. We thus observe a hierarchy of utility states: *Littered partly due to me, littered more than usual, littered the usual amount, clean*. For most citizens, the benefit of exiting the lowest state makes it worthwhile to bend over, pick up, and find a trash can. But they would not make that effort to transit upward from other states. The same underlying motivation, we believe, motivates many citizens’ thoughts about climate change: Their normal altruism for the future is reinforced by their guilt in contributing to the problem of greenhouse gases.

Consider another example, where concern about an act of commission plays a major role. It is a variant of the famous trolley problem proposed by Foot (1967):

A trolley is running out of control down a track, on course to kill five people. You can push a fat man in front of the trolley. He will be killed, but the trolley will be stopped and the five will be saved. Would you push the man?

Few people would push the man to his death.⁴ (Informal surveys suggest that economists are more willing to do so.) Commission makes a big difference, but it is not a trump card. If the numbers were 20 and 1, there would be more pushers. The same factors would influence policymaking for posterity. Actions that avoid imposing harm will get extra weight, but the magnitudes of costs and benefits will surely matter as well.

The implication of these examples is that in doing policy analysis for the very long run we should give greater weight to damage done to those alive in the distant future as a consequence of our actions than to damage caused by external events. While the issue of what is the appropriate reference point arises—is it not emitting any greenhouse gases? not increasing emissions? business as usual?—the implica-

⁴ The original version of the problem has a switch that can shuttle the trolley to an alternate track, where one person will be killed. This variant was proposed by Thomson (1985); she also notes that more individuals are willing to throw the switch, which seems to represent much lesser involvement than pushing someone.

tion would seem to be that a higher weight should be attached to damages caused by our emissions than to changes in consumption arising from other causes. Moreover, there would seem to be no obvious reason why this “caused damage” penalty should be felt less strongly as the length of time involved increases.

Summary The considerations adduced here all make a case for giving more weight to the distant future, taking more actions on its behalf, than would be implied by choosing “reasonable” parameters for the fundamental discounting equation. They also point in the direction of applying lower discount rates at longer horizons—a conclusion urged by others, notably Weitzman (1999) and Gollier (2008), on grounds related to uncertainty.⁵ We leave to future research the question of how the appropriate magnitudes can best be gauged.

2 Policy for posterity given uncertainty

Analysts who see the need for dramatic action to address the threat of climate change frequently invoke the tremendous uncertainty about the impact of climate change as an argument for more sweeping measures. They are dissatisfied with the conclusions that come out of standard economic models that employ discount rates of several percent. Weitzman (2008a and b) has attracted the most attention in this regard. In the previous section we identified a number of routes to serious action that did not involve uncertainty. Here we suggest that the effects of uncertainty are less clear cut than is often supposed. We take up two issues: the proper modeling of low probability catastrophic events and the impact of uncertainty and learning on the timing of mitigation policies. We then offer a brief observation on the implications of our analysis for geoengineering policy.

Valuing low-probability catastrophic losses due to climate change To simplify and aid intuition, let us leave the intergenerational context, and merely ask how the present generation should behave if it alone were empowered to save or potentially save the far future. That is, in the language of torts, we are the *last clear cost avoider*. Only if we “swerve” will the “accident”—in this case a climate catastrophe—be avoided. The simple situation is that we can invest today to reduce the risk of catastrophe to a generation living say 100 years into the future, and its successors. The first question, of course, is how large is the risk, and how much can we reduce it at what cost?

The second question is how we should value such a reduction. We address the second question, leaving the first primarily to the scientists. To address it, we must decide at the outset whether to rely on how most people appear to think, i.e., the behavioral approach, as opposed to what rational decision theory would prescribe. The answers can be far different.

⁵ Weitzman (1999, p. 29) recommends a discount rate of between 3% to 4% for 25 years, 2% for the next 50 years, 1% up to year 300, and then 0%. Gollier (2008) recommends a discount rate drifting down from 3.5% for the near term to 1% for a millennium hence. While these numbers may seem large, they are relevant in light of the EPA’s 10,000 year analysis of Yucca Mountain nuclear waste storage, which was recently exceeded by the newly announced 1 million year time horizon (see the introduction to this issue.)

The behavioral approach, as typified by Prospect Theory (Kahneman and Tversky 1979), would tend not to make a strong case for acting to reduce the risk of low probability catastrophes. This is because empirical evidence suggests that as long as there remains some risk of catastrophe people do not value very highly reductions in its magnitude from say 10% to say 1% very highly and certainly not as much as standard decision analysis would suggest they should. As Lichtenstein et al. (1978) points out, people overestimate the probability of low frequency lethal events, and underestimate the probability of more common fatal occurrences.

Climate change is a sufficiently serious problem, the dollars to be expended on it are sufficiently great, that we believe society should follow rational decision theory rather than common—indeed almost universal—behavioral propensities. We leave to our elected leaders how to work out the politics for pursuing the rational approach.

The rational approach merely requires two simple functions, one that gives the probabilities of various outcomes in response to our actions, and a second that values those outcomes. There is a great range of outcomes, to be sure. Within the category of catastrophe, there are those that merely represent grave loss of value, say equivalent to a reduction in GNP by 75%, and those that are truly cataclysmic, such that human life is no longer possible. To illustrate our approach, we will assume that there can be just two outcomes, normal and cataclysm. In effect, all catastrophes are upgraded to cataclysms. This conservative formulation is merited given that we shall conclude that even though our generation is the last clear cost avoider, and even though a cataclysm has a finite probability, we should not sacrifice unlimitedly, or indeed excessively, to avoid the cataclysm.

Our approach is inspired by the recent papers of Weitzman (2008a, b). He makes a convincing case that we cannot rule out a cataclysmic future outcome from climate change, in his terminology, that the “probability density function of climate catastrophes has an extreme tail that is heavy with probability” (2008b, p. 2). Following on this analysis, Weitzman foresees “potentially unlimited downside exposure.” This leads him to reject the traditional cost–benefit approach, within which discounting plays a major role.

Our formulation in the previous section allows for a departure from standard discounting models. Indeed, we are willing to attach reasonable values to costs going to far-future generations. Nevertheless, we draw a different conclusion than Weitzman. That is because whereas his valuation functions allow for unlimited downside exposure, i.e., utility going to negative infinity, ours do not.

Climate-change cataclysms in decisions made by society, we believe, should be treated on the same basis as personal cataclysms in personal decisions. Consider a bachelor with no dependents. He has a 1% chance of dying in the next week. That is as bad an outcome as he can imagine. This risk can be eliminated completely, for a price. What is the maximum amount he should pay? This question should be addressed with a von Neumann–Morgenstern utility function with two arguments: wealth, w , and a dummy variable representing life (1) or death (0). Basically, they solve the equation:

$$0.99U(w, 1) + 0.01U(w, 0) = U((1 - x)w, 1),$$

where x indicates the maximum fraction of wealth—which incorporates human capital—they would sacrifice. Most reasonable individuals, we believe, would give up much more

than 1% of their wealth. But we also believe that most would give up much less than 90% of their wealth. Basically, we are arguing that life at full wealth is given a utility value of 1, and death with full wealth is assigned a utility value of 0, that life with 10% wealth registers less than 0.99. That hardly seems unreasonable.

By way of contrast if we had assumed that death had infinite negative utility or alternatively that there were an unlimited range of terrible outcomes whose probability declined less rapidly than their disutility increased, we would have found that the individual was willing to make unlimited sacrifices to reduce the risk of cataclysm. Essentially this is what all of Weitzman's examples do. If one for example assumes that there is no outcome possible in which the utility losses are greater than those associated with a 99% drop in consumption, his strong conclusions no longer follow.

The argument could be made that even for an individual judging for himself, the extinction of human life is considerably worse than his own death, particularly if the extinction of life includes his death as well. To simplify thinking, leave aside the time dimension, and consider a world under a 1% threat of a current cataclysm that would kill all then alive and thus human life forever. Given current living standards in developed nations, what percent of its income would it give up to remove this threat? We doubt the median answer would be above 50%, and certainly not above 90%. This answer, whatever it is, is an upper bound on what should be paid to avoid such a threat in the future. We recognize, of course, that most of those who push for strong action against climate change would take exception to this hypothetical on one or both of two grounds: (1) The risk of cataclysmic outcomes given our current course with greenhouse emissions is way above 1%, perhaps above 90%.⁶ (2) The costs of dealing adequately with controlling greenhouse gas emissions should only be a small percent of income.

It is possible that there is a case based on low probability cataclysms for precipitous policy action, for far more than standard analyses suggest is appropriate. But this case we believe has to be made in a context that avoids approaches to infinite disutility.

Uncertainty and learning and the timing of climate-change mitigation All parties to the climate-change debate agree that it is afflicted with massive uncertainties. Interestingly, the two sides—those favoring greater and those favoring lesser immediate action—claim that the presence of dramatic uncertainties bolsters their argument. The former make the familiar point that the worse things could be, the more current action is necessary. The latter observe that with larger uncertainties, the benefit of waiting and learning what level of action is required is enhanced.⁷ As in

⁶ This begs the question of how to define a cataclysm. Few observers think global warming has much potential to wipe out human life forever, for example.

⁷ Both sides also claim irreversibilities as a possible source of support. Gollier et al. (2000) observes that scientific progress tends to enhance the efficiency of actions taken in the future, which tends to promote the learn-then-act principle, which favors an initially smaller amount of emissions reductions. Gollier and Treich (2003) address when and when not to delay action from a quite different perspective, namely the risk aversion properties of the utility function as represented by its second, third, and fourth derivatives. We look at the attractiveness of the learn-then-act approach in our quantitative model below, though the potential motivation for waiting within it is quite different from those just mentioned.

many a folk tale, both sides are right. The dispassionate economist would put it differently. Given two competing considerations, it is usually the case that for some parameter values one effect predominates and for others the other.⁸ We thus posit:

Hypothesis 1 *If there is uncertainty about damages, and that uncertainty will be resolved in the moderate future, then whether greater uncertainty will increase or decrease the fraction of total effort exerted in the first period will depend on parameter values.*

To prove this hypothesis, two simple numerical examples would suffice. The highly simplified model below is richer. It lays out the considerations that dictate when uncertainty should accelerate action and when it should delay it. There are two periods. They might be thought of as 10- or 20-year periods. Environmental costs are experienced in the second period due to first-period emissions. At the end of the second period, there is a tally of all future environmental damages due to first- and second-period emissions, what we will call *forward damages*.

Damages are uncertain, being either high or low. No learning about damages takes place during a period. At the end of the first period, the expected damages from emissions are learned, which will inform the level of reduction in second-period emissions.

The control variables are the levels of reduction in first- and second-period emissions, namely r_1 and r_2 . Without reductions, emissions will be 100 in either period. Thus, the r values can be thought of as the percentage of emissions reduced in a period. Reductions entail a cost, c . The cost function for either period is $c(r) = \gamma r^a$, where $a > 1$, implying increasing marginal costs, and γ is a constant. The uncertainty of concern is the level of damage due to climate change. The damage parameter, d , is either High or Low, indicated by the values H and L , with equal probability. The mean value of the damage parameter is always 1, implying that $L = 2 - H$. We assume that $1 \leq H \leq 2$. Damages from period 2 forward equal d times the sum of emissions in the two periods, namely $[d(200 - r_1 - r_2)]^q$, where $q \geq 1$.

There are also damages in period 2, the result solely of period 1 emissions. Period 2 damages are the fraction b multiplied by emissions to date multiplied by the damage parameter, where $0 \leq b < 1$. Thus, period 2 damages are $b[d(100 - r_1)]^q$. Overall, the objective function is:

$$\begin{aligned} \text{Minimize } & \gamma r_1^a + 0.5(\gamma r_{2H}^a + \gamma r_{2L}^a) + b[d(100 - r_1)]^q \\ & + 0.5([H(200 - r_1 - r_{2H})]^q + [L(200 - r_1 - r_{2L})]^q). \end{aligned}$$

The objective is to select r_1 at the outset and r_2 after d is learned so as to minimize the expected sum of the costs of reduction plus the costs of what damages turn out to be. For the moment, we leave discounting, risk aversion and period 2 damages aside, and set $q = 1$, to simplify exposition and intuition. Thus, this objective is consistent with either a framework to maximize the discounted sum of period expected utilities, or one that merely tallies total costs.

⁸ The structure of the model can matter as well, but we do not capitalize on this possibility.

The specific function we are considering is thus:

$$\text{Minimize } \gamma r_1^a + 0.5(\gamma r_{2H}^a + \gamma r_{2L}^a) + 0.5[H(200 - r_1 - r_{2H}) + L(200 - r_1 - r_{2L})].$$

The first term represents first-period cost; the second term is expected second-period cost, with a different value of r_2 selected depending on whether damages turn out to be High or Low. The third term represents the expected forward damages. The 0.5 values in the objective function represent the common likelihoods of the High and Low outcomes.

Our goal is to determine whether more uncertain damages lead to a larger or smaller fraction of total emissions reduction to be undertaken in period 1. Uncertainty is represented by the spread between H and L , which is fully captured by the value of H . We will investigate this question for different values of a , the cost parameter, and H , the damage parameter. Our concern is the effect of increased uncertainty about damages on period 1 action. First, we shall provide intuition for what we expect to find. Second, we shall produce a graph showing how $R_1 = r_1^*/(r_1^* + E[r_2^*])$ varies as a function of H for various values of a . The slope of any one of these curves gives dR_1/dH .

Note, that as a increases, cost becomes a more important consideration. To eliminate this factor, we can compute a different γ for each a to normalize outcomes. In our numerical calculations, we did this. Each γ was selected so that when there is no uncertainty, i.e., when $H = 1$, the optimal values of r are $r_1^* = r_2^* = 50$. That is, when damages are certain it is optimal to reduce emissions by 50% in each period.

Absent uncertainty and discounting, given that marginal costs are increasing, reductions should be equal in the two periods. Uncertainty on damages substantially complicates matters. First, relative to the certainty case, you are at least as well off. You could simply reduce 50 in each period and get the same expected outcome you did before. But you can do better by simply altering period 2's reductions depending on whether H or L applies. And you can improve on that outcome by adjusting period 1 reductions. How optimal period 1 reductions compare to overall optimal reductions is our subject of interest.

For the real world climate-change problem, we might consider ourselves to be in period 1. The question then is whether greater uncertainty on the magnitude of damages, holding expected damages fixed, should lead us to greater or lesser immediate action. The critical argument for lesser action is that we can wait, and then decide what to do. The critical argument for greater action is that if damages are uncertain and they turn out to be High, significant reductions will be required. Given the increasing marginal costs of reductions, if we wait till period 2, this will be very costly; thus, first-period action should be enhanced.

To illustrate, we wish to evaluate the difference in r_1^* between the situation where $H = 1$ and where $H = 1.5$, assuming that $q = 1$. Consideration of two extreme cases provides the intuition for our result. Let $a = 1.01$, so marginal costs only increase very slowly. Then there is little cost in waiting, with a rush of reductions if damage proves High. Thus, R_1 will be below 0.50 for this case. The optimal values turn out to be $r_1^* = 50$, $r_{2H}^* = 100$, and $r_{2L}^* = \epsilon$. Thus $R_1 = 50/(50 + .5(100 + \epsilon)) < .50$. Now let $a = 4$, so that marginal costs increase very rapidly. Rather than have second-period reductions have a 0.5 chance to be quite big, it is better to do a larger share of

reductions in the first period. Thus, R_1 will be above 0.50 for this case. The optimal values turn out to be $r_1^* = 50$, $r_{2H}^* = 57.236$ and $r_{2L}^* = 39.685$. Thus $R_1 = 50/(50 + .05(57.236 + 39.685)) = 0.508 > 0.50$.

The case where $a = 2$ represents a dividing line, since marginal costs increase linearly, and raising first-period expenditure by a unit to cut both High and Low expenditures by a unit neither saves nor costs money. Thus, if $a = 2$, the values of R_1 will be constant at 0.5, as $r_1^* = E[r_2]$ for all possible values of γ , H , and q .

For the case where $q > 1$, there are increasing marginal costs of damages. This makes it more important to have the ultimate reductions higher when H results. Once again, we adjust γ to assure that $r_1^* = 50$ when $H = 1$. Let $q = 2$. Consider as before the comparison between R_1 when $H = 1$ and when $H = 1.5$. If $a = 1.01$, R_1 goes from 0.50 to $30.397/(30.397 + 0.5(100 + 0)) = 0.378$. However, if $a = 4$, R_1 increases from 0.50 to 0.521. Note that once again, the a value dictates which side of 0.50 that R_1 will fall.

The graphs below shows the values of R_1 as H increases for various values of a , with $q = 1$. As is easily seen, dR_1/dH is negative when $a < 2$ and is positive when $a > 2$. This result holds independent of γ .

The value of a that minimizes R_1 for all values of H is about 3.7. The curves that cross are on opposite sides of this value, which explains the unusual appearance.

The next graphs are drawn for the case where $q = 2$, which could be thought of as some form of increasing marginal cost of damages. In comparison to the linear case ($q = 1$), it becomes more important when $q = 2$ to respond to High damages. Not surprisingly, r_1^* is a smaller fraction of total expected reductions when $a < 2$, and a greater fraction when $a > 2$. That is because when $a > 2$ it is much more costly to reduce emissions strongly in the second period should damages turn out to be high. In sum, the more rapidly the marginal cost of damages increases, the closer first-period expenditures are pulled to half of expected total expenditures.

These graphs provide the examples that prove Hypothesis 1. That is, and fully in accord with intuition, for different parameter values as uncertainty in damages increases, that may lead us to undertake a smaller or greater fraction of total expected reductions in emissions in period 1. The Appendix provides Fig. 3 as a topographical rendering of Fig. 2.

Inserting discounting into our analysis would not fundamentally change the results. A greater discount rate would lower the value of forward damages relative to second-period damages and costs, and lower second-period consequences relative to first-period costs. To illustrate, a constant less than 1 would be multiplied by second-period costs, making the first two terms in the objective function be $\gamma r_1^a + 0.5j(\gamma r_{2H}^a + \gamma r_{2L}^a)$ for $j < 1$. In effect, thus earlier costs and damages become more important. This would drive down r_1^* , but would not change the qualitative properties of dR_1/dH . The effect of period 2 damages is shown in Fig. 4 in the Appendix.

Risk aversion Risk aversion presents a more subtle story.⁹ Given risk aversion, it matters whether utilities are computed period by period, or all at the end. Following tradition,

⁹ Normally it is assumed that the government itself should not be risk averse given its ability to diversify, even across generations, but given the potential costs associated with climate change (e.g., curbing emissions, damages, adaptation, and geoengineering), risk aversion would seem merited.

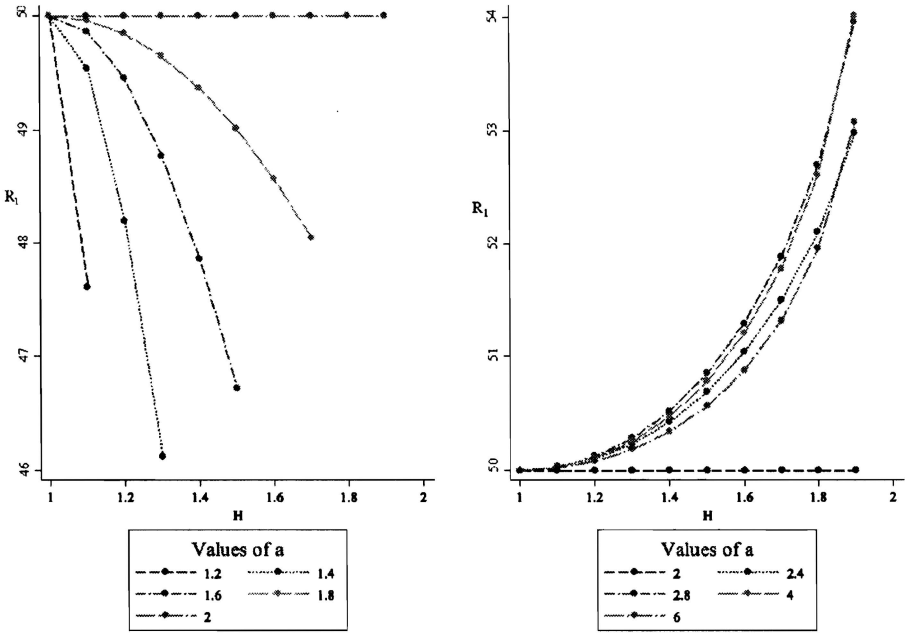


Fig. 1 Period 1 reductions as a percent of total reductions, $q=1$


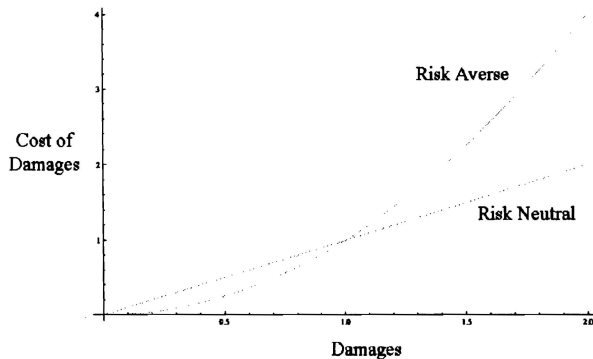


Fig. 2 Period 1 reductions as a percent of total reductions, $q=2$

we shall assume the former. This tradition makes particular sense in this context, since generations can hardly shuttle monies back and forth or make contingent trades between themselves. It is also critical to what variables risk aversion applies.

First consider risk aversion over costs of reduction. Remember, having $a > 1$ is not due to risk aversion, rather to escalating real costs of reducing emissions in a period. A tractable and plausible formulation for the utility function would have the disutility of costs rising like a power function. Thus, the loss from reductions in period i would be $[c(r_i)]^e$, where $e > 1$. This transforms the cost function to r_i^s , where $s = a^e$. For this utility function on costs, all our previous results go through if we just substitute s for a . Thus, there will still be a range of a where greater uncertainty about damages decreases the fraction of total reductions done in the first period ($a^e < 2$), followed by a range where it increases them ($a^e > 2$). Note that the cutoff would now be at a lower value of a , namely when $a^e = 2$ for any q .

Whatever the risk aversion on costs, let us now switch from risk neutrality to risk aversion on damages. This would mean that the cost of damages increases throughout, but more so for high damages (one could have thought of $q > 1$ as representing risk aversion, with expected monetary cost of damages held fixed as H varied). Risk aversion, in effect, just makes any distribution on positive expected damages worse, and the effect is greater the more variability there is in damages. Say that risk aversion was represented by a power function, where losses were raised to the power f . Then, rather than having damages raised to the q power, they would be raised to the $t = q^f$ power. Let us standardize with γ so that optimal reductions are 50 when $H = 1$, the central level of damages. Then the curve representing effective damages will switch from being a line through this point (risk neutrality) to being a convex curve through this point. These two curves are shown for risk aversion positing that mean damages are 1 and the cost of damages is squared.



Given the convexity of the cost of damages curve, and given that mean damages are 1, whenever there is variability in damages it effectively raises their cost. This implies that risk aversion on damages will lead to more first-period curbing when $H > 1$.¹⁰ Given risk aversion, greater uncertainty on damages will have two effects on first-period curbing. We explored the first effect in Hypothesis 1; it is ambiguous in sign.

¹⁰ If risk aversion is represented by a power function, its presence in effect raises the value of q , parallel to the way risk aversion on costs in effect raised the value of a .

The second effect comes because damages have become more consequential; it is unambiguously positive. Hypothesis 1 will still apply, but its negative-potential-effect component will apply over a more limited range.

This analysis highlights a number of factors that bear on the effects of increased uncertainty. A richer multi-period model would permit, we suspect, demonstration of the proposition that the speed with which uncertainty is resolved is likely to affect its impact on policy choices. Uncertainty that will be resolved quickly is likely to lead to more delay in taking costly steps than uncertainty that will only be resolved slowly.

The examples that we used in this section are obviously highly stylized but we believe they are sufficient to call into question the presumption that uncertainty about the impact of global warming necessarily calls for bearing more short run costs to provide for long term benefits than would be appropriate in its absence.

Geoengineering Most discussions of controlling climate change, including our own, have focused on curbing greenhouse gas emissions (mitigation). On a conceptual basis, however, we must also be aware that innovative measures could be taken that seek to lower temperatures by quite different means, or that seek to reduce damages from climate change. The three most prominent such measures are carbon sequestration, such as capturing carbon dioxide at the power plant level and pumping it underground, or planting trees to remove CO₂ from the atmosphere; adaptation, such as building sea walls to cope with sea-level rise or shifting agricultural land use; and geoengineering. It is notable that there has been a major shift towards more positive thinking about both sequestration and adaptation in the past fifteen years. Prominent climate scientists now study sequestration possibilities, a once verboten subject. The 2007 Bali Accord gave prominent mention to adaptation and to mitigation possibilities apart from greenhouse gas reduction, subjects almost absent from the Kyoto Protocols.

Geoengineering remains an approach more appealing to economists than other professionals addressing climate change. Perhaps the most discussed geoengineering measure would reduce the net amount of incoming solar radiation by shooting aerosols—sulfates being a frequent example—or solar reflectors (mirrors) into space.

Two prime concerns are announced about such measures: (1) They are unproven. To some critics, this includes the question as to whether they would be too costly even if they worked. (2) They will be upsetting the natural course, and could lead to worse environmental problems. Two other more politically attuned concerns may be less expressed, but surely influence attitudes: (3) The more hopes we hold about technological fixes, the less vigorous will be our efforts in curbing greenhouse gases, the original cause of climate change. (4) Failing to undertake geoengineering, would it have been successful (something we might never know) would be an error of omission. Undertaking geoengineering when it would prove detrimental would be an error of commission. The latter counts much more. A final concern is that geoengineering measures might prove attractive for individual nations—since cooling benefits are concentrated at the local level—even if they were net detrimental to the world as a whole.

We express no strong opinions on these matters, but make three observations that follow from the analysis in this and the previous section. (a) Even if geoengineering is a very unattractive option it may still affect the policy calculus if it can be called on quickly at some future juncture. Imagine that there is a 1% chance that global

climate change will have costs equivalent to 90% of world consumption, but that there is a geoengineering option that can avoid its consequences at a cost of 20% of world consumption. With low probability then geoengineering will look attractive in the future even if it does not today. As with the analysis in the first part of this section it serves to bound the cataclysmic possibilities and so affect current policy choices. (b) Considerations of learning become very important if geoengineering is even a possibility. If climate change does prove disastrous in any near-term scenario, geoengineering will be our only hope. Wigley (2006) observes: “More ambitious geoengineering, when combined with mitigation, could even lead to the stabilization of global mean temperature at near present levels and prevent future sea-level rise to a rate much less than observed over the twentieth century: aspects of future change that are virtually impossible to achieve through mitigation alone.” Curbing greenhouse gases, even severely, cannot alter concentrations significantly in a mere matter of years. If geoengineering is realistic even if highly costly and can be triggered in time then policymakers can treat a 20% loss as the worst possibility. (c) Arguably those most worried about global climate change should be most focused on geoengineering at least on an experimental basis. In fact, greenhouse alarmists tend to be strong opponents of geoengineering experiments, in part for reason (4) above, but also because they tend to believe that we have little understanding of all the ways human actions can seriously disturb the environment.

3 Reaction function issues

In any realistic policy analysis it is important to recognize that those whom the policymaker does not control are likely to adjust their behavior in response to changes in policy. Realistic policymaking must take this into account. Thus greater efforts to redistribute income are likely to lead to reduced work incentives. Public efforts to provide for retirement security, health coverage or help to the poor are to at least some degree likely to crowd out private efforts. On the other hand, some kinds of policies may have their effects magnified by the responses they generate. Think of demonstration projects or of efforts to make neighborhoods safe so more people walk on the street making them still safer.

In context of policymaking for the very long run these issues are especially important for the vast majority of what happens in the very long run is surely not under the control of current policymakers. And there is the further point that responses to any policies pursued are likely for good or for ill to grow over time. We consider in turn four aspects of the problem—the effects of well intentioned policies towards one social objective on other social objectives, the effects of the behavior of the current generation on the behavior of future generations, the effects of the behavior of one country where what is sought is the provision of an international public good, and the effects of policy on the rate of innovation.

Effects on alternative investments We have adduced in the context primarily of discussing global climate change a number of considerations arguing for giving greater weight to distant generations. While we have so far treated this as an argument for more aggressive climate policy, this conclusion may not follow. Should

we conclude that the future should be weighted more highly, there is the question of how to help the future. If we are going to spend \$100 billion to help people living 100 years from now for altruistic reasons, we should find the way to spend that money that yields the greatest benefit to them. Maybe we should be dramatically stepping up basic R&D efforts in a variety of areas, focusing on projects that may not yield fruit for decades. Perhaps we should be supporting a flotilla of young artists, some of whom might produce great works to be admired as we admire Van Gogh today. If these efforts would yield substantially more benefits to the citizens of 2210 than would enhanced efforts to curb climate change, then at least our intended beneficiaries would be better off (however, such efforts are directed, we would of course tally benefits along the path to 2210).

The critical point is that there are many projects that look attractive once we change our way of thinking about the future. In principle all such projects should compete against one another, and we should start with the ones that are most favorable. In practice, decisions are typically made about particular opportunities. Thus the standard response of global climate change policy advocates to suggestions that alternative policies might yield even greater long-run benefits is to say that “maybe and if so we should do them as well.” Perhaps, but at a certain point the capacity to allocate resources to the distant future is stretched and indeed if enough resources were so allocated levels of consumption would decline which would operate to raise the discount rate applied to future benefits by increasing the expected growth rate.

Probably the right approach in addressing very long-term investments, analogous to the one taken by Feldstein (1964) in thinking about discounting standard public investments, would assess the impacts they have on other very long-term investments, and would then apply a shadow price to value these impacts. Thus a global climate initiative that is for reasons economic or political likely to crowd out basic scientific research should be judged more harshly than an equivalent project that would affect only consumption.

The intergenerational game In 2008, the G-8 nations agreed to cut emissions in half by 2050. Even if we care about the future, and are willing to sacrifice for it, we would prefer to have others make that sacrifice for us. Thus, people around in 1998 but no longer alive, who cared about 2050 emissions, should be pleased to be free riders on the altruism of their successors, namely us. The flip side of this argument is that we would hardly be pleased to take costly action from 2010 to 2020 only to find that people in 2020 then said they could then slacken their effort in response.

If one used a standard reaction function analysis, the more we do today, the less will be done tomorrow. That recognition would dampen enthusiasm for action, and could lead to an unfortunate slackers’ equilibrium. This argument is explored in Arrow (1999) and in the original classic paper by Phelps and Pollak (1968).

But appropriately framed, the direction of influence could go in the opposite direction: One good deed could promote another. The current generation’s significant effort could set a precedent for generosity. Say its successor framed its choice issue as: “We have a good policy going. Do I want to continue, or do I want to break the chain?” Then, each successive generation might find it optimal to continue. In effect, even though each generation would sell the furniture if it thought

its successor might do so, none would sell, since each would be confident that all successive generations would see matters the same way.¹¹

Think of a more home-grown example. You have inherited valuable antique furniture from your grandparents. You are now providing for your own descendants. One possibility is to sell the antiques, and invest the monies. You conclude that even taking their sentimental values for the furniture into account, they would be better off with the money. Moreover, you do not like taking care of valuable furniture. But you would feel some guilt about selling the furniture that has come down through the generations. If you would feel sufficient guilt, you should keep the furniture. Moreover, if you were willing to sacrifice a certain amount for your descendants, you are better off than having sold the furniture. You would be better off as well if you merely had indifference curves between their welfare and your welfare. You should save some more for them in addition.

This formulation leads to an additional question on intergenerational values. We have discussed altruism toward a specific future generation. If we have such altruism, so will generations in between. In theory, we should value their altruistic benefits as well. Thus, when generation 1 saves for generation 3, it should count the benefit that gives to generation 2. In theory, this process could also go backward in time. Generation 2, when providing for 3, must consider generation 1's concerns. Conceivably backward altruism should extend even to generations no longer present. If those generations could not have counted on such altruism, they would have been less generous themselves.¹²

The picture we have tried to present is of a procession of generations all with concerns for the future. The challenge to them is to make provisions for that future in the most efficient manner, given that the generations cannot readily contract with one another. In these grand tacit intergenerational games, framing matters greatly. Actions will be favored if they can be framed in such a way as to call forth increased effort from future generations directed at the very long run, rather than offsetting reductions in effort.

Here Schelling's (1957) work on focal points is helpful. Actions framed as noble precedents, and established around bright-line markers, are likely to stimulate positive responses. By contrast, complex or overly ambitious schedules for action, which are likely to get short-changed a bit even by us, provide both strong temptations for greater chiseling in the future and a negative future response to our positive actions now.

International considerations As noted in Section 1, the longer the time horizon the more cosmopolitan the appropriate perspective for policymakers in any one nation. There remain however several questions about policymaking for the very long run, particularly in a context where decisions are made possibly through international negotiation at the national level.

¹¹ Stochastic shocks to the system, say a generation might bounce to *poor* and have to sell the furniture would complicate the problem, and would make it much harder to maintain the precedent.

¹² Their total altruism toward a generation incorporates both its felicities, e.g., benefits from consumption, and its altruistic concerns. The authors consider backward as well as forward altruism (Zeckhauser and Fels 1968).

First, there is the aspect of voluntary provision of an international public good. Absent some form of binding agreement a country considering taking costly actions to reduce greenhouse emissions has to recognize that unless its choices are affecting the choices of other countries they are likely to yield only modest benefits to its citizens. Thus individual countries, even the largest emitters, have strong incentives to ride free. Cheap riding, e.g., making small reductions, may be slightly preferred, to provide a warm glow to or diminish the guilt of one's environmentally oriented citizens, or to defuse foreign claims that one is free riding. Absent an effective international agreement, how should a nation think about whether to ride free, ride cheaply, or pitch in fully in producing the good?

Second, if one assumes that international negotiation is possible, several issues immediately arise. Can agreements be enforced in a world of sovereign nations? In the absence of enforcement mechanisms will nations adhere to their commitments? And will inevitable concerns about these questions lead agreements to unravel so strongly that even effective enforcement mechanisms will be useless? Assuming away these problems, if different nations think about posterity differently how should their preferences be aggregated? What is the right form of an agreement that is expected to influence behavior over a very long time period? On the one hand it is important to be flexible in a highly uncertain world. On the other, it is necessary to make firm commitments in areas where only very long-term investments can produce the gains needed. There is an inevitable tension between establishing enduring focal points for action, and responding effectively to new information.

Third, how should nations set policy prior to an agreement being reached, or when it is anticipated that an existing agreement will be modified? There is the possibility that more extensive efforts will encourage reciprocity on the part of others. But there is also the possibility of a ratchet—that whatever progress has been made will be taken as given, so the reference point for judging anti-emission efforts will become that much more challenging.

All of these issues become much more difficult when the putative international agreement involves the very long run. For this means that precedents set may matter greatly for good or ill, it means that urgency may be diminished, and it means that enforcement issues become more profound when agreements will still be important long after those who negotiate them are no longer involved. Their successors may not feel bound, particularly if their fundamental values differ from those who signed the accord, say because their nation underwent a transformation from socialism or communism to capitalism.

Technology Policymaking for posterity encounters an additional critical reaction function issue that arises from the impact of policy commitments on innovation. Two seemingly contradictory bits of folklore, each distilled from experience, come to mind. On the one hand, it is generally believed that cost estimates for major policy changes, like Medicare, or putting a man on the moon, or buying a new weapons system, or completing the Big Dig always err massively on the low side. And it is generally supposed that the greater a step into the unknown a policy change represents, the greater the underestimate will be. On the other hand, there is substantial experience in the environmental area suggesting that when clear commitments are made, the cost of ultimately meeting them is well below initial

estimates. Examples often cited include the US sulfur oxide program, the initial smog controls in Los Angeles, and the ban on CFCs.

Implicitly these two pieces of folklore represent arguments about how projected future costs should be discounted. Proponents of either argument believe that it applies with special force in the case of climate change, in part because of the very long time periods involved. Which of these lessons of experience carries more weight for climate change policy? Most important, do current estimates under- or overweight the costs of controlling emissions? The issue is quite important as extrapolation from the former set of examples could easily suggest that commitments as large as those involved in achieving huge reductions in carbon emissions could easily cost five times or ten times initial estimates, whereas the environmental examples can be read as suggesting that initial estimates should be cut by two-thirds or more.

Our suspicion is that judging which consideration is more weighty in any given policy context depends on evaluating the source of the estimates. Estimates for new government programs typically are heavily influenced by those most knowledgeable about them, who in turn are likely to be enthusiastic and optimistic. On the other hand, industries facing the prospect of being regulated have every incentive to produce high estimates of the costs of meeting any given goal.

4 Conclusion

How vigorous should our current efforts be to combat climate change? More generally how much should we do for posterity? This essay has clustered the answer under the four Ds: discounting, disaster, distinction, and decision analysis. All four factors must be considered in arriving at an answer. The mnemonic is appropriate for a 4D world, with time the fourth dimension.

Discounting When looking at a problem that stretches across many generations, the economist's standard approach of exponential discounting may be problematic. We raised the possibility based purely on preferences—one often asserted by non-economists—that standard formulas are an inappropriate guide for long-distance discounting. That is, we may, or should, value a generation 300 years hence no different than one 100 years hence, if say we are saving it from our comet collision.¹³ We are agnostic on some such deep discounting questions, but believe that if current curbing efforts merit vigorous action, it requires that we have a discounting framework that makes the future weighty.¹⁴ In the realm of discounting, we observed that all generations will not be created equal. If we expect them, as we do, to have more members and to be richer, the first clearly calls for a higher weight. The second almost certainly does for what we have labeled amenity goods.

¹³ We reiterate the concern that having no diminution in weighting through time, as might come from an independent survival risk, implies that a small benefit going to all future generations has infinite value.

¹⁴ As we observed in the introduction, any standard discounting approach using traditional discount rates would give far-future generations virtually no weight.

Disaster Costly efforts to combat climate change are only potentially worthwhile if future disaster threatens. Weitzman (2008a, b) provides a compelling argument why we should expect the distribution of costs from climate change to have a fat tail, implying high expected costs at some future date. That finding alone, however, does not tell us how energetic an effort to pursue now. In addition, a utility function must be attached to the possible outcomes. Here we part ways with Weitzman. Though following normal utility theory analysis, we do not allow for infinitely negative outcomes. Thus, unlike him we will accept a finite probability of a complete disaster before sacrificing nearly everything. Nonetheless, we accept the notion that any consideration of climate-change policy must start with an assessment of the magnitudes and likelihoods of potential disasters.

Distinction There are many ways to provide for the far future. If one employs a normal utilitarian framework, where the relevant arguments are the utilities of the different generations (incorporating altruistic components), then our goal should be to find the policy that provides the greatest benefit to the far future per unit cost to us.¹⁵ With that formulation, it would seem that climate-change policy would have to compete with other measures that help the far future, such as medical research. Say the choice were between devoting an incremental 2% of our GNP on a continuing basis to curbing greenhouse gases or to enhancing medical research, perhaps evenly divided for the moment between cancer, Alzheimer's and tropical diseases.¹⁶ Which would accomplish more for the far future on an expected value basis? Say that study revealed that it was medical research. We still might prefer to curb the gases, if due to guilty altruism, or the difference between acts of omission and commission, the utility cost to us was less. Such a what-do-the-donors value argument is familiar in the discussion of in-kind transfers, where the middle class appears to prefer medical care for the poor as opposed to cash transfers. Society appears to draw strong *distinctions* between different types of assistance. Affirmative action surely gets an extra push because of our history of slavery, for example. If we find it cheaper to in effect avoid littering the atmosphere than making the same expenditures to promote medical knowledge, that calculus should inform our climate-change policy.

Decision analysis Curbing greenhouse gases in an appropriate manner is a long-term dynamic optimization process. The greater are uncertainties, the greater is the risk of waiting. Even if variability does not increase expected damages, it raises the urgency to curb now because reduction costs increase at an increasing rate. But there is a countervailing consideration. Increased uncertainty also raises the potential for learning along the way, which could reduce optimal current efforts. Our apparently simple and stylized model showed that when these two effects tug in opposite

¹⁵ We do not focus attention on any single generation, but think of the far future as a stream of future generations.

¹⁶ The NIH 2008 budget was \$29.5 billion. 2% of GNP would produce roughly ten times that amount. Presumably an increase in expenditures would take place predominantly in the for-profit sector, say achieved through some tax incentive, but still it would take a few years to ratchet up spending significantly without gross waste.

directions, greater uncertainty—assuming that it will be at least partially resolved in the intermediate future—may lead to more or less vigorous current efforts.

We believe that these four Ds represent puzzle pieces that must be put together if we are to formulate appropriate climate-change policies. Economists and their fellow travelers have a comparative advantage for guidance on *discounting* and *decision analysis*. Scientists have deep knowledge on disasters, and can provide an assist on decision analysis. Philosophers and psychologists can supply insights on discounting and distinction.

Our principal conclusion is that analyses that focus on merely one or two of the four Ds will be no more than a two-dimensional rendering. Policymaking for posterity must be formulated for a four-dimensional world.

Acknowledgments We thank Guan Yang for his research assistance with the literature, and Nils Wernerfelt for both research assistance and the computation of the results of our uncertainty and learning model. Kip Viscusi provided helpful comments.

Appendix

Topographical renderings of first-period reductions and the effect of second-period damages

In both Figs. 1 and 2 on the graphs on the right, as a increases, the value of R_1 first increases, then decreases. Below is another graph for $q = 2$, showing R_1 along the vertical axis as a function of values for H and a , given that $b = 0$, $\gamma = 3$, and $q = 2$. Figure 3, consistent with the right side of Fig. 2, shows that R_1 is maximized for some value of a around 4. R_1 decreases as a shrinks to 2 or rises to infinity. The graph is qualitatively similar for $q = 1$.

We set aside second-period damages thus far; that is set we $b = 0$. Obviously making $b > 0$ will increase R_1 for any value of H . The next question is how do positive second-period damages affect dR_1/dH . This depends upon the parameter values in question. For example, in Fig. 4 ($\gamma = 4$, $a = 1.5$, and $q = 2$), for higher values of H , the value of b matters more; for lower values of H , b matters very little.

Fig. 3 Topographical rendering of Fig. 2





Fig. 4 First period reductions given second period damages



Thus, given a large amount of uncertainty, the marginal effect of second-period damages on dR_1/dH is greater, albeit decreasing.

References

- Arrow, K. (1999). Discounting, morality, and gaming. In P. R. Portney, & J. P. Weyant (Eds.), *Discounting and intergenerational equity*. Washington, DC: Resources for the Future.
- Arrow, K., Solow, R., Portney, P. R., Leamer, E. E., Radner, R., Schuman, H. (1993). Report of the NOAA Panel on Contingent Valuation." <http://www.darrp.noaa.gov/library/pdf/cvblue.pdf>.
- Cowen, T. (2008). Social discount rate. *The New Palgrave Dictionary of Economics Online*, 2nd ed.
- Cropper, M., Aydede, S. K., & Portney, P. R. (1994). Preferences for life saving programs: How the public discounts time and age. *Journal of Risk and Uncertainty*, 8(3), 243–246.
- Feldstein, M. (1964). The social time preference discount rate in cost benefit analysis. *Economic Journal*, 74(294), 360–379.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5–15. Reprinted in Bonnie Steinbock and Alastair Norcross, eds. *Killing and letting die*, 2nd ed. (New York: Fordham University Press, 1994), pp. 266–279.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40, 351–401.
- Gollier, C. (2008). Discounting with fat-tailed economic growth. *Journal of Risk and Uncertainty*, 37 (2/3).
- Gollier, C., & Treich, N. (2003). Decision-making under scientific uncertainty: The economics of the precautionary principle. *Journal of Risk and Uncertainty*, 27, 77–103.
- Gollier, C., Jullien, B., & Treich, N. (2000). Scientific progress and irreversibility: An economic interpretation of the 'precautionary principle'. *Journal of Public Economics*, 75, 229–253.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology*, 4(6), 551–578.
- Nordhaus, W. D. (2007). A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45(3), 686–702.
- Phelps, E. S., & Pollak, R. (1968). On second-best national saving and game—Equilibrium growth. *Review of Economic Studies*, 35(2), 185–199.
- Portney, P.R., and Weyant, J.P. (Eds.) (1999). *Discounting and intergenerational equity*. Washington, DC: Resources for the Future.
- Schelling, T. (1957). Bargaining, communication, and limited war. *Conflict Resolution*, 1(1), 19–36.
- Stern, N. (2007). *The economics of climate change: The Stern Review*. Cambridge, UK: Cambridge University Press.
- Thaler, R. H., & Benartzi, S. (2001). "Save more tomorrow: Using behavioral economics to increase employee saving." <http://faculty.chicagogsb.edu/richard.thaler/research/SMarT14.pdf>.
- Thomson, J. J. (1985). The trolley problem. *Yale Law Journal*, 94(6), 1395–1415.

- Weitzman, M. L. (1999). Just keep discounting, but.... In P. R. Portney, & J. P. Weyant (Eds.), *Discounting and intergenerational equity*. Washington, DC: Resources for the Future.
- Weitzman, M. L. (2007). A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45(3), 703–724.
- Weitzman, M. L. (2008a). *Some dynamic economic consequences of the sensitivity inference dilemma*. Mimeo: Harvard University.
- Weitzman, M. L. (2008b). *On modeling and interpreting the economics of catastrophic climate change*. Mimeo: Harvard University.
- Wigley, T. M. L. (2006). A combined mitigation/geoengineering approach to climate stabilization. *Science* (314), 20 October 2006, 452–454.
- Zeckhauser, R., & Fels, S. (1968). Discounting for proximity with perfect and total altruism. Discussion Paper Number 50, Harvard Institute of Economic Research.