

Karp , Larry S.

Working Paper

Sacrifice, discounting and climate policy : five questions

CESifo Working Paper, No. 2761

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Karp , Larry S. (2009) : Sacrifice, discounting and climate policy : five questions, CESifo Working Paper, No. 2761, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/30615>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Sacrifice, Discounting and Climate Policy: Five Questions

Larry Karp

CESIFO WORKING PAPER NO. 2761
CATEGORY 10: ENERGY AND CLIMATE ECONOMICS
AUGUST 2009

PAPER PRESENTED AT CESIFO VENICE SUMMER INSTITUTE, JULY 2009

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

Sacrifice, Discounting and Climate Policy: Five Questions

Abstract

I offer a selective review of discounting and climate policy. Analytic and numerical models show that different assumptions greatly change the degree to which decisions about climate policy depend on the discount rate. I discuss a claim that standard models exaggerate the current generation's sacrifices needed to internalize climate damages. This claim, if correct, affects the role of discounting. I argue that the assertion that the risk of catastrophic damage overwhelms discounting is unfounded. I show that the claim that we "view the world in perspective" implies hyperbolic rather than constant discounting.

JEL Code: C61, C73, D63, D99, Q54.

Keywords: climate change, discounting, intergenerational conflict, catastrophic risk, hyperbolic discounting.

Karry Karp
Department of Agricultural and Resource Economics
207 Giannini Hall
University of California
USA – Berkeley CA 94720
karp@berkeley.edu

July 30 2009

Paper prepared for CESifo Venice Summer Institute workshop on "The Economics and Politics of Climate Change".

This paper benefited from numerous conversations with Tomoki Fujii, Rolf Golembek, Yacov Tsur, and Christian Traeger and comments from participants of the CESifo Venice Workshop; the usual disclaimer applies. I thank Penelope West for constructing an example that convinced me that Proposition 1 is correct, and I especially thank Niklas Mattson for showing me the crux of the proof of that Proposition.

1 Introduction

Protecting future generations from climate change requires that we alter this generation's investment and consumption decisions. The discount rate translates future utility or consumption into the same units as current utility or consumption, thereby making inter-generational comparisons possible. The discount rate therefore can be important in making recommendations about current climate policy. The discount rates for utility (the pure rate of time preference) and for consumption (the social discount rate) are different, but related, objects. I recognize this difference where it is important, but for much of the discussion I merely use the term "discount rate". I provide a (very selective) overview of several recent discount-related topics in climate policy, including the magnitude of effects, the extent of intergenerational conflict, the degree to which catastrophic risk overwhelms other considerations, and our perspective on the distant future. In addition to providing a summary for economists who have not previously encountered one or more of the issues that I raise, I hope that a unified discussion of the issues will be a useful synthesis.

First, I consider the evidence on the importance of the discount rate in evaluating climate policy (Section 2). The discount rate is especially likely to be important if protecting future generations from climate change requires that the current generation make sacrifices, not merely reallocate investment across different uses. A somewhat heterodox view is that although reallocation of investment across different uses in the current period must occur, sacrifices might not be necessary; in any case, these sacrifices have been exaggerated by previous researchers because of a modeling error. To the extent that this view is correct, the discount rate becomes less important in determining climate change policy – or at least the focus changes (Section 3).

Section 4 considers the relation between discounting and the desire to avoid catastrophes. Discount rates and catastrophe-avoidance are logically distinct topics, but recent papers claim that the risk of catastrophe swamps any consideration of discounting. I am skeptical of this claim. Section 5 takes up the question of how to treat the distant future in models of the environment. I provide a new result showing that spatial perspective is analogous to *hyperbolic* discounting, not merely discounting. To the extent that our picture of the world involves perspective over both space and time, this result provides another motivation for using hyperbolic discounting in climate policy models. The final section discusses recent modeling developments that

offer useful new approaches for thinking about the climate problem.

2 How important is the discount rate?

Most economists working on climate policy think that protecting future generations from climate-related damages will require sacrifices by the current generation, typically in the form of reduced consumption. The hypothesis that reducing emissions requires a reduction in current consumption, and the recognition that climate-related damages will likely arise decades or even centuries in the future, suggests that the discount rate is an important determinant of optimal climate policy.

The comparison between current costs and future benefits in simple settings provides much of our intuition about the sensitivity, to the discount rate, of optimal climate policy. However, apparently similar settings can result in much different policy levels and policy sensitivity to the discount rate. A simple example shows that similar models can have quite different implications about discounting. After this example, I point out that more complex climate change models can also lead to different views of the importance of discounting.

In the simple example, suppose that by reducing the flow of utility by x units per period we can avoid a loss of 100 units of the flow of utility over (T, ∞) . If T is known, we would be willing to give up the flow $x = e^{-\rho T}100$ over $(0, \infty)$ when the pure rate of time preference (PRTP) is ρ . If $T = 200$, x changes by a factor of 55, ranging from 13.5 to 0.25, as ρ ranges from 0.01 to 0.03 (1% to 3% per annum).

In the variation of this example, suppose that instead of being a known parameter, T is an exponentially distributed random event with hazard rate h . In this case, the maximum reduction in the flow of utility, beginning today, that we would accept in order to eliminate the hazard under risk neutrality and a constant PRTP of ρ is

$$x' = \frac{h}{h + \rho}100. \tag{1}$$

In order to make this variation comparable with the original example, suppose that the expected time-to-occurrence is $\frac{1}{h} = T$, so

$$x' = \frac{1}{1 + \rho T}100.$$

For $T = 200$ as before, the maximum amount that we would be willing to pay, x' , changes by a factor of 2.3 (compared to 55 in the original formulation), ranging from 33.3 to 14.3 as ρ ranges from 0.01 to 0.03. The ratio of willingness-to-pay in the two examples is

$$\frac{x'}{x} = \frac{\exp(\rho T)}{1 + \rho T} > 0.$$

Both the level of the willingness to pay and the sensitivity of the level with respect to the discount rate, are quite different in these two examples, despite their superficial similarity.

With intuition based primarily on the first version of the example, where climate policy is very sensitive to the discount rate, many economists focused on the Stern Review's (2006) (hereafter SR) discounting assumptions (although others emphasized different aspects of the report, including assumptions about abatement costs and climate related damages). SR chose a PRTP of $\rho = 0.001$, an elasticity of marginal utility of $\eta = 1$, and a growth rate of 0.013, implying a social discount rate (SDR) of $r = 0.014$ (or 1.4%).

Nordhaus (2007) illustrates the importance of these discounting assumptions by comparing the results of three runs of the DICE model. Two of these runs use combinations of the PRTP and η consistent with a SDR of about 5.5%. With these values, the optimal carbon tax in the near term is approximately \$35/ton Carbon (or \$9.5/ ton CO₂), and the optimal level of abatement in the near term about 14% of Business as Usual (BAU) emissions. A third run, using the SR's values of the PRTP and elasticity (together with the DICE assumptions about growth) led to a carbon tax of \$350/ton and a 53% level of abatement, close to the level that the SR recommends. Thus, the carbon tax increases by a factor of 10 and abatement increases by a factor of $\frac{53}{14} = 3.8$ with the decrease in the SDR. Nordhaus also describes experiments by others that illustrate the sensitivity of optimal climate policy to discounting assumptions.

In a different context (focused on the effect of catastrophic damages) Nordhaus (2009) compares optimal policy under a PRTP of 0.015 and 0.001, holding other DICE parameters (including the elasticity of marginal utility) at their baseline levels. He reports that the reduction in PRTP increases the optimal carbon tax from \$42/tC to \$102/tC. This 2.4-fold increase in the optimal tax is much less than the 10-fold increase reported in Nordhaus (2007), where both the PRTP and the elasticity are changed.

Discount rate	.01	.03	.05	.07
Abatement first period (%)	25.1	9.8	5.2	3.3
Abatement 10th period (%)	18.4	8.8	5.2	3.5

Table 1: Percent reduction in emissions from BAU level

Paraphrasing, the increase in the tax from \$42/tC to \$102/tC leads to a fall in per capital income (during the period when it is lowest, presumably the first period) from \$6,801 to \$6,799, i.e. about 0.03%. In view of parameter uncertainty, a 0.03% reduction in per capita income, equivalently a \$15 billion increase in aggregate abatement costs (at Gross World Product – GWP – of \$50 trillion), is close to rounding error. If it is correct that using a \$102/tC tax rather than a \$42/tC tax would reduce current GWP by about \$15 billion, then the criticism of the SR recommendations is that they harm future generations, not the current generation. This may be the right interpretation of the criticism, but it is not the one usually applied to climate policies that are deemed “radical”. Section 3 takes up this issue.

Karp (2005) provides a different perspective, using a linear-quadratic model, calibrated to reflect abatement costs and climate-related damages that are of the same order or magnitude as in DICE. Table 1 shows the optimal percentage reductions in BAU emissions in the first and the 10th period (100 years in the future) for different discount rates. In this stationary, partial equilibrium model, there is no growth, so the pure rate of time preference and the social discount rate are the same. However, Nordhaus’s experiments, described above, reduce the pure rate of time preference from 1.5% to essentially 0, leading to a reduction in the social discount rate of approximately 1.5%. Those results are therefore (roughly) comparable to the results in Table 1, where the discount rate varies by a magnitude of 2% across columns.

The lower discount rate leads to a substantially higher, but not extraordinarily higher level of abatement. For example, the decrease in the discount rate from 3% to 1% increases abatement by a factor of 2.5, close to the proportional increase in the optimal tax in Nordhaus (2009) when the PRTP falls from 1.5% to 0.1%. In this linear-quadratic model, decreasing the discount rate from 3% to 1% has approximately the same effect on the near-term optimal policy as does tripling the level of damages corresponding to any level of GHGs.

Fuji and Karp (2008) provide a more involved analysis of the role of dis-

counting, still within the context of a model with a single state variable. The SR reports values of the “output gap” resulting from climate-related damages, in eight different periods and corresponding to three different abatement scenarios. Using these values, we calibrated a scalar model in which the state variable, P_t , represents the fractional loss in income due to climate damages and the reallocation of past resources away from standard investment and toward climate-related measures. Our objective was to create a tractable model that is consistent with the orders of magnitude of costs and benefits in the model underlying the SR. The control variable, x_t , is the fraction of GWP devoted to avoiding climate change, so the actual amount of consumption, as a fraction of the amount that would have been available in the absence of climate damage and abatement effort, is $(1 - P_t)(1 - x_t)$. The consumption loss due to mitigation expenditures and remaining climate-related damage, as a fraction of the no-damage no-control level is $\Delta_t = 1 - (1 - P_t)(1 - x_t)$. The model allows a constant growth rate in income and a constant elasticity utility of consumption. We used a constant annual growth rate of 1.3%, as in the SR.

Table 2 shows optimal steady state values of the state and control variables, P and x , and the aggregate damage measure, Δ , for four combinations of the PRTP and elasticity of marginal utility, η . With a growth rate of 1.3%, these values encompass SDRs ranging from 1.31% to 5.6%. Because all of the steady state values are small in absolute value, the percentage change rather than absolute change in values is the relevant statistic. For example, with $\eta = 1$, increasing the PRTP from 0.1% to 3% increases the fraction of steady state expenditures (relative to available income) by about 19%. Increasing both the PRTP and η , so that the social discount rate increases from 1.31% to 5.6%, increases steady state expenditures x by 26%. These numbers refer to steady state levels, whereas the numbers discussed previously consider the abatement and tax levels in the current period or “middle future”. However, for the calibration in Fujii and Karp (2007), the steady state values are more sensitive to the SDR than are interim values.

I noted above that in the Nordhaus (2007) experiments, abatement increases by a factor of 3.8 when the SDR falls from 5.5% to 1.4%. Because abatement costs are convex, expenditures will increase by a larger factor. In other words, the 26% (steady state) increase in Fujii and Karp (and lower interim sensitivity) appears to much lower than the increase that Nordhaus finds.

Figure 1, the graph of the steady state costs, Δ , as a function of steady

	$\eta = 1$			$\eta = 2$		
PRTP	P^*	x^*	Δ	P^*	x^*	Δ
0.1%	0.0137	0.0084	0.0220	0.0147	0.0076	0.0222
3.0%	0.0165	0.0068	0.0232	0.0185	0.0062	0.0246

Table 2: Steady state damages and expenditures

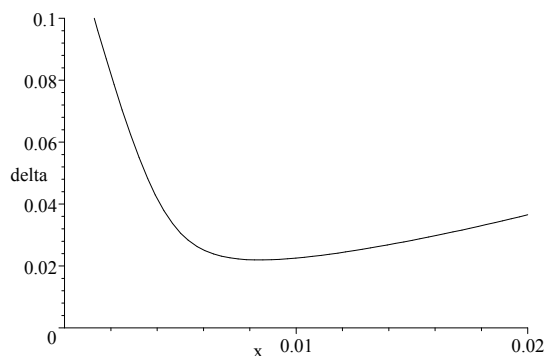


Figure 1: The graph of steady state costs, Δ , as a function of steady state expenditures, x

state expenditures x , provides a partial explanation for this relative lack of sensitivity. This graph reaches a global minimum at $x = 0.00845$ (the optimal level under zero discounting) and falls rapidly for smaller values of x . The optimal x under the SDR 1.31% nearly achieves the global minimum. The characteristics of Δ mean that small increases in expenditure, below the global optimum, achieve significant reductions in total costs. Therefore, a very low SDR achieves nearly the global minimum, and even substantially larger SDRs take us close to the global minimum. Because initial expenditures (compared to steady state expenditures) are even less sensitive to the SDR, the entire trajectory is “relatively insensitive” to the discount rate.

Arrow (2007) examines the effect of the discount rate on our willingness to avoid climate change, posing the question in terms of growth rates rather than levels of damages. Although this way of presenting the tradeoff is not standard, given the state of knowledge about both climate-related damages and abatement costs, it does not seem obviously inferior to the familiar alternative. Arrow considers the case where climate policy requires a sacrifice of

1% of the flow of GWP in perpetuity. The policy eliminates climate-related damages, leading to growth in GWP at the rate 1.4% (without damages) instead of 1.3% (with damages). For a constant elasticity utility function with $\eta = 2$, he notes that we would be willing to adopt the policy, and sacrifice the 1% of GWP, provided that the PRTP is less than 8.7%. This level is much higher than those used in all integrated assessment models, implying that it is an easy choice to adopt this policy.

In Arrow's experiment, GWP is higher under the climate policy at every point in time after the first ten years. That is, the climate policy requires a modest sacrifice for a short period of time, in exchange for what eventually becomes a large increase in consumption. Perhaps this seems like a reasonable description of climate policy to many people.

I conclude that the answer to the question posed in the title of this Section is ambiguous. Certainly the discount rate matters, and in some settings it matters hugely – as our intuition suggests. In other settings it appears to matter less than we might have thought. Evidence of the sort summarized here cannot be conclusive, because it all comes from specific models or specific ways of presenting the tradeoff between abatement costs and avoided damages.

3 Does climate policy require sacrifice?

There are two kinds of reasons that the discount rate can be important for climate policy. Different types of investments lead to different profiles of future benefits, the comparison of which depends on the discount rate. Therefore, the allocation of a given level of investment, across different investment opportunities, depends on the discount rate. Also, as emphasized in the previous section, the discount rate affects the willingness of the current generation to make sacrifices, i.e. investments, for future generations. A widely held view, and the source of much of the controversy about discount rates, is that taking into account the climate externality requires that the current generation reduce its current consumption, imposing a cost on that generation. Foley (2007) explains why this view may be incorrect, and Rezaei, Foley and Taylor (2009) (hereafter RFT) explain why the mistake may be of practical significance.

Consider the frontier between current and future consumption in the absence of abatement opportunities, shown as the solid curve in Figure 2. (For

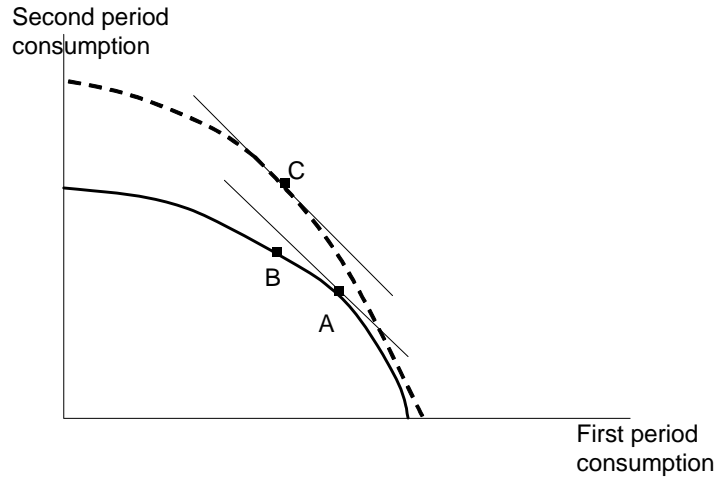


Figure 2: Point A: the constrained optimum with abatement set equal to 0. Point B: the BAU equilibrium. Point C: the first best optimum with abatement.

the time being, ignore the letters in this figure.) Internalization of the climate externality opens people’s eyes to the possibility of abating, in order to leave future generations with a cleaner environment. If society does not internalize these damages, the only way for the current generation (the “first period”) to leave a bequest for future generations (the “second period”) is to leave them a higher stock of man-made capital. By internalizing future damages, the current generation can also benefit future generations by leaving them a larger environmental stock. Therefore, internalization, together with opportunities for abatement, shifts out the consumption possibility frontier, making it possible to increase consumption in both the current period and the future (the dashed curve in Figure 2). Since climate policy makes it possible to increase both current and future consumption, the extent of the intergenerational conflict may be less and the focus of the discount rate different than widely thought.

The above point is important and widely understood, although sometimes neglected in policy discussions. Nordhaus (2007), page 695, considers this possibility. He discusses a “fiscal experiment” in which society reduces

emissions and in addition¹

... undertake[s] fiscal tax and transfer policies to maintain the baseline consumption levels for the present (say fifty years). The optimum might have slightly lower consumption in the early years, so the fiscal-policy experiment would involve both abatement and fiscal deficits and debt accumulation for some time, followed by fiscal surpluses and debt repayment later. In essence, this alternative keeps consumption the same for the present but rearranges societal investments away from conventional capital (structure, equipment, education and the like) to investments in abatement of greenhouse gas emissions (in “climate capital”, so to speak).

...The reason why the [SR’s] approach is inefficient is that it invests too much in low-yield abatement strategies too early. After fifty years, conventional capital is much reduced, while “climate capital” is only slightly increased. The efficient strategy has more investment in conventional capital at the beginning, and can use those additional resources to invest heavily in climate capital later on.

The second paragraph quoted makes it appear that Nordhaus’s (2007) rejection of the SR recommendations is based on a comparison of the returns to different types of investment. However, the emphasis on the discount rate in his paper, and comments about the current generation making sacrifices for possibly wealthier future generations, shifts the focus to intergenerational transfers. Regardless of his intention, much of the discussion about climate policy amongst economists, climate scientists, and the general public, does emphasize the cost to the current generation of internalizing climate damage.

Integrated assessment models (IAMs) can measure the cost to the current generation of optimal climate policy by comparing current consumption under the optimal program with consumption under a baseline that does not internalize climate change. RFT claim that some of these IAMs choose the wrong baseline, systematically exaggerating the baseline level of current consumption and thereby upwardly biasing the estimate of the consumption

¹Nordhaus’s mention of fiscal deficits and debt accumulation is puzzling in this context. The reallocation of investment away from man-made and toward natural capital is consistent with a decreased, increased, or unchanged debt.

cost to the current generation that would arise from internalizing climate damages. Based on this result, they conclude that the true cost to the current generation (in units of foregone current consumption) of optimal policy may be negligible.

It is worth emphasizing that the choice of the baseline used in IAMs does not affect the validity of the models' policy recommendations. The choice of the baseline obviously does affect the estimated present discounted stream of benefits of optimal policy – the difference between optimal and baseline levels; the baseline also affects the comparison of optimal and Business as Usual (BAU) current consumption, and therefore affects the magnitude of the proposed sacrifice to the current generation.

In order to understand this claim, consider a simple model with two state variables, (man-made) capital stock and the stock of GHGs (natural capital). Production creates emissions, leading to higher GHG stocks and higher environmental damages. Society can reduce current emissions by allocating some current production to abatement, but this requires a reduction in either current consumption or investment in man-made capital. More complicated models, in which society can reduce future abatement costs by investing in clean technology, do not change the intuition, so there is no advantage here in considering such models.

The first best optimum can be found by solving a standard optimization problem in which a single decision-maker chooses abatement, consumption, and investment in each period, subject to technology (the production function) and the equations of motion for capital and the GHG stock.

How should we calculate the baseline? It seems reasonable to take the baseline as the competitive equilibrium in which decision-makers have rational expectations but do not internalize the environmental damages. In this BAU baseline, agents do not pretend that environmental damages do not exist; however, they act as if their decisions have no effect on these damages. Finding this rational expectations competitive equilibrium requires finding a trajectory of investment and consumption decisions and a corresponding sequence of GHG stocks and environmental damages, such that the investment/consumption trajectory is a competitive equilibrium taking as given this sequence of environmental damages. This equilibrium seems like the most reasonable baseline; I refer to it as the BAU baseline. Finding the BAU baseline requires solving an equilibrium problem, not a simple optimization problem.

RFT claim that some prominent IAMs do not calculate this baseline cor-

rectly, and that instead they merely solve an optimization problem in which abatement is constrained to be 0 (or some other fixed, non-optimal level); call this the “constrained optimization problem”. That is, RFT claim that these IAMs merely “turn off” the abatement control, and solve the problem in which a decisionmaker chooses investment and consumption optimally, respecting the production function and the equation of motion for capital and GHG stocks and understanding the relation between environmental stocks and damages. This problem is a standard optimization problem, not an equilibrium problem, but it does not have any obvious significance. It results in a rational expectations equilibrium, because the resulting future trajectory of damages equals the trajectory that the decisionmaker anticipates when choosing current consumption and investment. However, if agents really treat climate damages as an externality, then they should not take into account how current investment alters future emissions, thereby affecting future damages.

An alternative assumes that agents ignore future damages. That alternative is at least consistent with the assumption that agents treat damages as an externality. However, it does not lead to a rational expectations equilibrium, because agents are systematically wrong about the level of damages. In IAMs where the climate affects productivity, this error causes agents to be systematically wrong about the return on capital.

Both the constrained optimization problem and the alternative in which decision-makers ignore damages are unsatisfactory baselines. The first turns off the abatement control and the second turns off environmental damages. The first leads to a rational expectations equilibrium but allows the decisionmaker who invests today to internalize *future* (although not current) damages, leading to an internally inconsistent model. The second uses an internally consistent model, in which agents ignore both current and future damages, but it does not produce a rational expectations equilibrium. Decision-makers in the first model are “too smart” and they are “too dumb” in the second alternative, at least according to current modeling norms.

In order to show why using the constrained optimization solution as a baseline leads to an exaggerated estimate of the reduction in current consumption required by optimal climate policy, I use a two period model. RFT use a genuinely dynamic model calibrated to be consistent with DICE, and they show the empirical importance of this result. My objective here is only to provide intuition about this claim, so I use a stripped down two-period

model.²

This model requires that I compress the dynamics in a way that appears artificial. In particular, I assume that the emissions resulting from current production contribute to the current GHG stocks, reducing current as well as future output. A more descriptive model would have current emissions caused by current production contribute to GHG stock and damages only in the next period, but that greater degree of realism would require a three period model.

When current capital stock is k and lagged GHG stock is S_{-1} , and there is no abatement, emissions are equal to $F(k)$, so the current GHG stock is $S_{-1} + F(k)$. By choice of units, output in the current period in the absence of climate-related damage is also $F(k)$. Actual output in the current period, taking into account climate damage, and excluding the possibility of abatement, equals

$$Y = D(S_{-1} + F(k)) F(k).$$

The function $D(S)$ represents environmental damages, with

$$D(S) \leq 1 \text{ and } D'(S) < 0.$$

An increased stock of GHGs increases damages and decreases output. In the first period, the current capital stock, k , and lagged GHG stocks, S_{-1} , are predetermined. By choice of units, set

$$D(S_{-1} + F(k)) = 1 = F(k). \tag{2}$$

Suppose that $F(0) = 0$; if the capital stock is 0, there is neither output nor emissions.

If society abates the fraction μ of emissions, total output is

$$Y = D(S_{-1} + (1 - \mu) F(k)) F(k)$$

²The other alternative baseline, in which the decision-maker simply ignores damages (the “too dumb” scenario), also leads to an incorrect estimate of the cost of internalizing climate change, but the direction of the bias is ambiguous. Ignoring future damages leads to an upwardly biased estimate of future income, thus lowering the incentive to invest in the current period. However, ignoring future damages also leads to an upwardly biased estimate of the marginal product of capital, thus increasing the incentive to invest in the current period. In any case, the “too dumb” scenario is probably not as important as the “too smart” scenario, because to my knowledge no IAM creator has been accused of using it.

and abatement costs are $Y\gamma(\mu)$, with $\gamma(0) = 0$. Remaining output is allocated between consumption and investment:

$$D(S_{-1} + (1 - \mu)F(k)) \cdot F(k) \cdot (1 - \gamma(\mu)) = c + I.$$

The stock of GHGs accumulates according to

$$S = F(k)(1 - \mu) + S_{-1},$$

and the stock of capital depreciates in a single period, so next period's capital equals the current period's investment:

$$k' = I. \tag{3}$$

The utility of current consumption is c and there is no discounting, so in period 0 the decisionmaker wants to maximize the sum of current and future consumption, $c + c'$. In the second (and last) period it is not optimal to invest anything, since capital has no value after that period.

Consider three possible equilibria.

(i) In the first best equilibrium, agents choose consumption, abatement, and investment in the first period, and consumption and abatement in the second period. They take into account the income constraints in the two periods and the equations of motion.

(ii) The constrained optimal equilibrium represents the case where agents are constrained to set abatement equal to 0, but they understand that investment in the current period increases damages in the second period, via the increase in second period emissions $F(k')$.

(iii) In the BAU equilibrium, agents take aggregate emission and the trajectory of damages as exogenous. Since they take emissions as exogenous, they have no reason to abate, so $\mu = 0$. But also, since they take emissions as exogenous, they do not take into account the fact that first period investment increases second period emissions and damages, a fact that reduces second period consumption. That is, in the BAU equilibrium agents overstate the marginal value of investment.

In both the BAU and the constrained optimal equilibrium $\mu = 0$ in both periods; this equality and equation (2) imply $c = Y(1 - \gamma(\mu)) - I = 1 - I$. In these two equilibria, first period emissions equal 1 and second period consumption equals

$$c_2 \equiv D\left(S_{-1} + 1 + \underline{F(I)}\right) F(I),$$

where the underlined term is the additional second period emissions due to first period investment. The maximand for the constrained optimal program is $1 - I + c_2$. I assume that c_2 is concave in I so that the first order condition is sufficient for a maximum. The first order condition for the constrained optimum is

$$1 = \left[\underline{D'(\cdot) F(I)} + D(\cdot) \right] F'(I), \quad (4)$$

where the argument of the damage function in the second period is

$$S = S_{-1} + F(k) + F(I) = S_{-1} + 1 + F(I). \quad (5)$$

The presence of the underlined term in this first order condition shows that in the constrained optimal equilibrium, the current generation internalizes future damages. They do not abate current emissions, but in choosing current investment they recognize that higher capital in the next period leads to increased damages.

In the BAU equilibrium, agents take next period damages as given, so the first order condition for investment is

$$1 = [D(\cdot)] F'(I).$$

The fact that

$$\left[\underline{D'(\cdot) F(\cdot)} \right] F'(I) < 0$$

and the second order conditions to the optimization problems imply that investment is lower, and therefore first period consumption is higher, in the constrained optimal equilibrium, relative to the BAU equilibrium.

Figure 2 shows the graph of c_2 as a function of c_1 when abatement is constrained to 0, the solid curve. The constrained optimum is point A and the BAU equilibrium is point B . As noted above, the constrained optimum involves higher first period consumption, in order to reduce future damage, compared to the BAU level. The dashed curve shows the graph of second

period consumption as a function of first period consumption, when the decisionmaker chooses abatement optimally. The with-abatement frontier lies outside the no-abatement feasible set.

In the first-best optimum, when the decisionmaker chooses both first period investment, I , and abatement in the two periods, μ and μ' , the maximand is (using $F(k) = 1$)

$$[Y \cdot (1 - \gamma(\mu)) - I] + [D(S_{-1} + 1 - \gamma(\mu) + (1 - \gamma(\mu')) F(I)) F(I)],$$

where the two bracketed terms are first and second period consumption. The first order condition for investment is

$$1 = [D'(\cdot) (1 - \gamma(\mu')) F(I) + D(\cdot)] F'(I), \quad (6)$$

where the argument of the damage function is now

$$S^{\text{opt}} = S_{-1} + 1 - \gamma(\mu^{\text{opt}}) + (1 - \gamma(\mu'^{\text{opt}})) F(I) \quad (7)$$

rather than the level in equation (5)

The right sides of both equations (4) and (6) are decreasing in I by the second order condition. In any reasonable equilibrium, S^{opt} is less than the level under the constrained optimum. If that relation holds and if damages are linear ($D'' \equiv 0$), then the graph (as a function of I) of the right side of (4) lies below the graph of the right side of (6). In this case, (using the fact that in general $\mu^{\text{opt}} > 0$) investment in both man-made and natural capital is higher in the first best compared to the constrained optimum.³ However, in view of the relation between investment in the constrained optimum and in BAU, investment in man-made capital under BAU can be higher than in the first best outcome. In that case, the optimal policy involves less investment in man-made capital and more investment in natural capital ($\mu > 0$) relative to BAU. Figure 2 shows a situation where first period consumption is approximately the same under BAU and under the first best, although clearly the relation is ambiguous in general.

RFT calibrate a model with natural and man-made capital, using abatement costs and environmental damages similar to those in DICE. A major point of their exercise is to show that using the constrained optimum instead of BAU as the baseline greatly exaggerates the cost to the current generation (in terms of foregone current consumption) of behaving optimally.

³This result could be overturned if utility were non-linear in consumption.

IAMs emphasize the effect of optimal policy on the change in the present discounted value of the stream of future utility, rather than the change in the *profile* of consumption (or welfare). For this reason, it might appear that RFT have set up a straw man. I disagree with that judgement for two reasons. First, there is a widespread perception that protecting future generations from climate change requires reductions in consumption in the near term. The perception may be wrong, or at least exaggerated, and it is worth demonstrating that this is so in a model that makes assumptions about damages and abatement costs similar to those used by prominent IAMs. Second, when IAMs use the constrained optimum rather than BAU as the baseline, they understate the increase in the present discounted flow of welfare arising from optimal climate policy. This point is also worth making clearly.

4 Do catastrophes swamp discounting?

Weitzman (2008) examines the effect of parameter uncertainty on the social discount rate. Using a two period model, representing the current period and the distant future, he calculates the marginal expected value of transferring the first unit of certain consumption from the present into an uncertain future. His chief result is open to several interpretations. In my view, a “modest” interpretation is correct and useful. A controversial interpretation is that the result undermines our ability to sensibly apply cost-benefit analysis to situations where there is uncertainty about catastrophic events. A “corollary” to this interpretation is that the recognition of catastrophic events makes discounting a second order issue. I think that both the controversial interpretation and the corollary to it are incorrect.

In order to explain these points, I consider a simplified version of his model. Let c be the known current consumption, c' the random future consumption, x the number of certain units of consumption transferred from the current period to the future, β the discount factor (representing the pure rate of time preference), and u the utility of consumption. The social discount factor for consumption, i.e. the marginal rate of substitution between “the first” additional certain unit of consumption today and in the future, is

$$\Gamma = -\beta E_{c'} \left(\frac{\frac{du(c'+x)}{dx}}{\frac{du(c-x)}{dx}} \right) \Big|_{x=0}. \quad (8)$$

The chief result is that, under the assumptions of the model, $\Gamma = \infty$; Weitzman dubs this result “the dismal theorem”.

The model includes a number of important features, including: (i) the uncertainty about c' is such that there is a “significant” probability that its realization is 0; (ii) the marginal utility of consumption at $c' = 0$ is infinite; and (iii) it is possible to transfer a certain unit of consumption into the future. Features (ii) and (iii) are assumptions, but (i) is a result of the model; most of the modeling effort goes to constructing a setting that yields this feature. Weitzman achieves this by assuming that the variance of c' is an unknown parameter, and the decisionmaker’s subjective distribution for this parameter has “fat tails”; roughly speaking, this means that the probability of extreme values of the variance does not approach 0 “too quickly”. As applied to the problem of climate change, all three features of the model, and especially (i) and (ii), have been criticized. Although there is tremendous uncertainty about the effects of climate change, some modelers feel that it is reasonable to bound the possible damages away from levels that would result in consumption at a level where the marginal utility is infinite, or even enormously large.

Those criticisms are important, but a more fundamental issue is how to interpret the dismal theorem. A modest interpretation is that uncertainty about the distribution of a random variable can significantly increase “overall uncertainty” about this random variable, leading to a much higher risk premium (and therefore a much higher willingness to transfer consumption from the present into the future) relative to the situation where the distribution of the random variable is known. This modest interpretation is not controversial.

An extreme interpretation is that under conditions where the dismal theorem holds, society should be willing to make essentially any sacrifice to transfer a unit of certain consumption into the future. That interpretation is also not controversial, because it is so obviously wrong. Even with 0 discounting (in this two period stationary model with the same utility function in both periods), we would never be willing to transfer to the future more than half of what we currently have.

The controversial interpretation is that the dismal theorem substantially undermines our ability to sensibly apply cost-benefit analysis to situations with “deep uncertainty” about catastrophic risks. The basis for this claim is that in order to use the social discount rate given in equation (8), we need to modify the model so that Γ is finite. Weitzman suggests ways of doing this,

such as truncating a distribution or changing an assumption about the utility function or its argument, in order to make Γ finite. The alleged problem is that the resulting Γ is extremely sensitive to the particular device that we use to render it finite. Because we do not have a consensus about how to achieve this finite value, we do not have a good way to select from the many extremely large and possibly very different social discount rates. In this setting, it is difficult to use cost-benefit analysis.

This controversial interpretation is not persuasive. Horowitz and Lange (2008) identify clearly the nub of the misunderstanding; I rephrase their explanation. Nordhaus (2008) also identifies this issue, and he provides numerical results using DICE to illustrate how cost-benefit analysis can be used even when damages are extremely large.⁴

The problem with the controversial interpretation is that the value of Γ in equation (8) is essentially irrelevant for cost benefit analysis. This expression, which is evaluated at $x = 0$, gives the value of the “first” marginal unit transferred. The fact that the derivative may be infinite does not, of course, imply that the value of transferring one (non-infinitesimal) unit of sure consumption is infinite. If we want to approximate the value of a function, it makes no sense to use a Taylor approximation evaluated where that function’s first derivative is infinite. We would make (essentially) this mistake if we were to use equation (8) as a basis for cost-benefit analysis with climate policy. The value of the derivative evaluated at $x = 0$ is almost irrelevant for cost-benefit analysis of climate policy. The only information that we obtain by learning that the derivative is infinite is that a non-infinitesimal policy response must be optimal. This fact is worth knowing, but it obviously does not create problems for using cost-benefit analysis.

Although I think that the controversial interpretation is unsound, it has a “corollary” that is not so easy to dismiss. This corollary states that catastrophic risks swamps the effect of the pure rate of time preference. The idea is that since the expectation of the term in parenthesis in equation (8) is so large, the magnitude of β , and thus of the pure rate of time preference, is relatively unimportant. Nordhaus (2009), despite his criticism of the

⁴These numerical results are interesting, but someone who accepts the controversial interpretation of the dismal theorem will not regard them as a convincing counter-argument to that interpretation. All of the numerical experiments arise in a deterministic context, and in that respect they do not really confront the dismal theorem, which makes sense only in a setting with risk and uncertainty (Weitzman 2009). For this reason, I think that Horowitz and Lange’s very simple treatment of the problem is particularly helpful.

controversial interpretation of the dismal theorem, endorses this view:

...discounting is a second-order issue in the context of catastrophic outcomes. ... If the future outlook is indeed catastrophic, that is understood, and policies are undertaken, the discount rate has little effect on the estimate of the social cost of carbon *or to the optimal mitigation policy*.

(Emphasis added.) Presumably this judgement derives from numerical experiments with DICE, but Nordhaus does not provide evidence to support his assertion; the only *reported* comparison involving discounting holds other DICE parameters at their baseline level, i.e. at levels not consistent with catastrophic damages.

This corollary may hold in specific settings, but it would be surprising if it is a general feature of catastrophic risk. The magnitude of the expectation of the term in parenthesis in equation (8), evaluated at $x = 0$, can certainly swamp the magnitude of β ; but I have just noted that the former term is irrelevant for cost-benefit analysis (beyond telling us that non-infinitesimal policy is optimal).

In order to get a sense of whether the corollary is likely to hold, and also to illustrate why the controversial interpretation of the dismal theorem is not persuasive, I use an example with $u = \frac{c^{1-\eta}}{1-\eta}$. Set $\beta = \exp(-\rho T)$ and choose a unit of time equal to a century. With this choice of units, ρ is the annual pure rate of time preference expressed as a percent. Suppose that c' takes the value c with probability $1 - p$ and the value 0 with probability p . For $p > 0$ the right side of equation (8) is infinite, as in the dismal theorem. The optimization problem is

$$\max_x (u(c - x) + \beta [(1 - p) u(c + x) + pu(x)]).$$

Normalize by setting $c = 1$, so that x equals the fraction of current consumption that we transfer into the future. With a bit of manipulation, the first order condition for the optimal x is

$$\rho = \frac{1}{T} \ln \left((1 - p) \left(\frac{1 - x}{1 + x} \right)^\eta + p \left(\frac{1 - x}{x} \right)^\eta \right). \quad (9)$$

Using $\eta = 2$ and $T = 1$ (so that the “future” is a century from now), Figure 3 shows the relation between the *annual* percentage pure rate of time

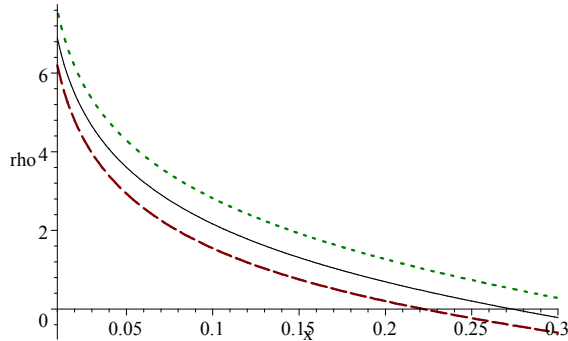


Figure 3: The relation between the transfer, x , and the annual percentage discount rate, ρ , for $p = 0.05$ (dashed), $p = 0.1$ (solid) and $p = 0.2$ (dotted), with $\eta = 2$.

preference, ρ , and the optimal value of x for p equal to 0.05, 0.1 and 0.2. The figure illustrates the obvious point that the fact that the expression in equation (8) is infinite does not cause any problem in determining an optimal value of the transfer. It also illustrates the not-so-obvious point that catastrophic risk does not swamp the effect of discounting, in determining the optimal level of the transfer. Changes in p have an effect on the optimal x that appear to be of the same order of magnitude as changes in ρ .

We get a bit more insight into the relative importance of the PRTP, ρ , and the probability of catastrophe, p , by taking ratio of the elasticities of the optimal x with respect to these variables. Define the elasticities and the ratio of elasticities as:

$$\frac{dx}{dp} \frac{p}{x} = \lambda; \quad \frac{dx}{d\rho} \frac{\rho}{x} = \alpha; \quad \tau = -\frac{\alpha}{\lambda}.$$

For $\tau > 1$, x is more sensitive to discounting than to the probability of catastrophe, and the reverse holds for $\tau < 1$. Tedious calculation shows that for the isoelastic example with $\eta = 2$

$$\tau = \frac{1}{p} \frac{\ln \frac{1}{x^2} \frac{(x-1)^2}{(x+1)^2} (x^2 + 2px + p)}{2x + 1} (x^2 + 2px + p),$$

where $x = x(\rho, p, T)$ is defined implicitly by equation (9). The ratio of elasticities depends explicitly on p and x and it depends implicitly on T , because that variable affects the optimal value of x .

Continuing the example, set $p = 0.05$: the probability that in one century consumption (absent the transfer) falls to 0, equals 0.05. Figure 4 shows the graph of τ for $x(\rho, p, T) \in [0.13, 0.21]$ with $\eta = 2$ and $p = 0.05$. With these parameter values (and $T = 1$ century), $x = 0.13$ is the optimal transfer to the future when the annualized PRTP is 1%, and $x = 0.21$ is the optimal transfer when the annualized PRTP is 0.1%.

There is not a general agreement about the “reasonable” level of the PRTP. I think that a 3% rate would be widely regarded as a high value, 1% as a moderate or low value, and 0.1% (the value used in the SR) as a very low value. Therefore, consideration of the interval $x(\rho, p, T) \in [0.13, 0.21]$ means that we are (implicitly) considering moderately low to very low PRTPs. Over this interval of x , τ ranges (inversely) over $[0.16, 1.27]$; in addition, $\tau \geq 1$ for $x \leq 0.15$, or equivalently $\rho \geq 0.74$. (It is worth reminding the reader that due to the choice of time units, ρ is the annualized PRTP expressed as a percent.)

In summary, for the example with $p = 0.05$, $\eta = 2$, and $T = 1$, the optimal transfer is more sensitive to the PRTP than to the probability of catastrophe ($\tau > 1$), whenever the PRTP exceeds an annualized level of 0.74%. Readers who think that PRTPs in excess of 0.74% include the range of reasonable rates, would conclude that for this example, the optimal transfer is more sensitive to the discount rate than to the probability of catastrophe.

The lowest value of τ in this experiment, corresponding to $\rho = 0.1$, is 0.16. Thus, for extremely low discount rates, the optimal transfer is $\frac{1}{.16} = 6.25$ times more sensitive to the probability of catastrophe than to the discount rate. As $\rho \rightarrow 0$, the optimal transfer $x \rightarrow 0.222$, and as $x \rightarrow 0.222$, $\tau \rightarrow 0$.

That is, for PRTPs arbitrarily close to 0, $\frac{1}{\tau}$ becomes arbitrarily large. Therefore, there is a situation (extremely low ρ) where consideration of the risk of catastrophe does swamp discounting. However, this situation is rather trivial, and we do not need a specific model, such as I have used here, to recognize it. In any two-period, or finite horizon setting, the optimal transfer from the present to the future is bounded above; for example, in the two period setting we obviously would never transfer more than half of current consumption, whatever the discount rate or the probability of catastrophe. For a given probability of catastrophe, the optimal transfer reaches its maximum at $\rho = 0$ (under the assumption that $\rho \geq 0$). Therefore,

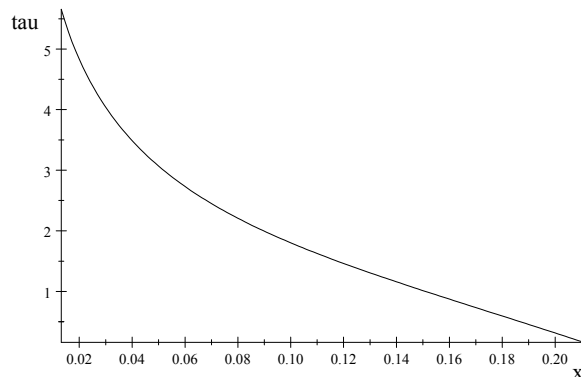


Figure 4: The ratio of elasticities, τ , as a function of x for $\eta = 2$ and $p = 0.05$.

for ρ extremely close to 0, the elasticity of the transfer with respect to ρ must approach 0, so τ must approach 0. In other words, the comparative statics question that this experiment addresses is interesting only for ρ bounded away from 0.

5 How do we view the distant future?

Here I discuss whether hyperbolic discounting provides a reasonable model of how people view the world, and then I discuss some implications of my answer. The major effects of today's GHG emissions might occur decades or even centuries in the future. Therefore, the evaluation of future welfare is likely to matter to climate policy. The Ramsey formula for the social discount rate adopts the perspective of an infinitely-lived agent who has a constant pure rate of time preference. The absence, in the real world, of financial markets that make it possible to directly transfer a unit of consumption from today into any point in the future, is not important. The agent today can indirectly transfer consumption into any point in the future simply by making a succession of one-period transfers. The absence of these long-term financial markets may matter if the agent has time-inconsistent preferences, such as arise with a non-constant, e.g. hyperbolic, pure rate of time preference.

The pure rate of time preference measures our willingness to take utility away from a future generation in order to increase the utility of an earlier

generation. This measure plausibly depends on our ability to make distinctions among different generations. Based on introspection, I think that we are better able to distinguish among people who are closer to us, either in space, time, or genetically, compared to people who are further from us. If this view is correct, then we discount hyperbolically, not at a constant rate.

Ramsey (1928) remarked “My picture of the world is drawn in perspective. ... I apply my perspective not merely to space but also to time.” Perspective applied to space means that objects further in the distance appear smaller. In this regard, perspective applied to space is analogous to discounting applied to time: events further in the future appear less significant. It is perhaps not obvious that perspective applied to space is analogous to *hyperbolic* discounting.

I use an extremely simple example to support this claim. Figure 5 is a drawing of a railroad in one-point perspective. The rails, which are parallel in reality, appear to converge at the horizon. The horizontal lines, the railroad ties, are actually evenly spaced, but more distant ties appear to be closer together. The letters $A, B, C...$ denote the successive ties, and also the apparent length of the ties. Although the actual lengths are equal, the apparent lengths diminish: $A > B > C...$. The letters $a, b, c...$ denote the apparent distance between ties. The person looking at this railroad is standing in front of the first tie, A . If this person were ubiquitous, she would correctly perceive the rails to be parallel and the ties to be evenly spaced. Since she is located in a particular position, the rails appear to converge and the ties to get closer together in the distance. The taller she is, the more her view resembles that of the floating deity.

There are two equivalent interpretations of hyperbolic discounting in this context. One interpretation is that the apparent rate of decrease of the distance between ties falls as the ties become further away from the person. This interpretation is analogous to the idea that the interval of time, say a year, seems smaller further in the future. For example, delaying an event for a year, starting today, may seem more significant than delaying an event from 500 years to 501 years from now. The second interpretation is that the rate of decrease of the apparent length of the ties falls as the ties become more distant. This interpretation is analogous to the idea that exchanging a dollar today for a dollar one year from now appears to involve a greater cost than exchanging a dollar 500 years from now for a dollar 501 years from now.

More formally:

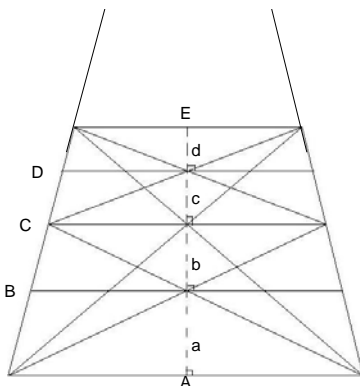


Figure 5: A railroad drawn in one-point perspective. The sides of the trapezoid (the rails) are parallel in reality, and the horizontal lines (the railroad ties) are evenly spaced in reality.

Definition 1. *A person with one-point perspective has hyperbolic discounting with respect to space if and only if*

$$(i) \frac{c}{b} > \frac{b}{a} \text{ and } (ii) \frac{C}{B} > \frac{B}{A}.$$

These two inequalities correspond to the first and the second interpretations described above. More precisely, inequalities such as these must apply for any three successive distances, not simply for the first three as shown. I express Definition 1 using the first three distances only to conserve notation.

Proposition 1. *A person with one-point perspective has hyperbolic discounting with respect to space.*

We do not give a second thought to our perception of space. If one accepts Ramsey's analogy between space and time, then the recognition that our spatial perspective is hyperbolic makes it easier to accept that our temporal perspective is also hyperbolic. I provide this proposition in order to help make hyperbolic discounting seem like an obvious, not an esoteric choice.

Does hyperbolic discounting matter? There are two steps to answering this question. I first discuss how hyperbolic discounting affects the way that we think about problems, and then I discuss how it affects the trade-offs we are willing to make for climate policy.

There are a number of ways that we might think of the equilibrium in a problem with hyperbolic discounting, but two choices stand out. First, we might consider the case where the decision-maker at time 0 can commit to future policies. This choice solves the time consistency problem by assuming it away. It strikes me as an unappealing choice in the case where climate policy unfolds over different generations. The current decision-maker has limited ability to influence her successor, except by affecting the state of the world that the successor inherits. The second alternative is to model the equilibrium as the outcome to a sequential game amongst a succession of policymakers. Different definitions of equilibrium exist. A particular choice, the Markov Perfect equilibrium, assumes that decisionmakers condition their actions on only the directly payoff relevant state variables, such as the stock of GHG and the stock of abatement capital.

A frequent question is whether the Markov Perfect equilibrium (set) is positive or normative. It shares features of both. Insofar as it is an outcome to a game, rather than a single-agent optimization problem (and insofar as we think that the model is descriptive of the real world), it is positive. In another respect it is normative, because models of this sort typically assume a single decisionmaker in each generation. That is, they assume that the intra-generational problem has been solved, in order to focus on a the inter-generational problems. A more genuinely positive model would include both intra- and inter-generational games; but the “positive” (or predictive) characteristic of such a model would still depend on the extent to which we think it is descriptive.

Turning to the second part of the question, in some but not all circumstances hyperbolic discounting (using a Markov Perfect equilibrium) has a significant effect on policy prescriptions. Section 2 notes that in some case the policy prescription is not very sensitive to the pure rate of time preference. This is the case in Fujii and Karp (2008), and there the distinction between constant and hyperbolic discounting is also not very important. In the linear-quadratic model (Karp, 2005) I found the optimal policy to be quite sensitive to the pure rate of time preference; there, hyperbolic discounting can be quite significant to policy. Karp and Tsur (2008) study a model with event uncertainty, where the hazard rate increases over time unless the decisionmaker incurs a cost to stabilize (but not eliminate) the hazard. Equation (1) shows how the hazard rate affects the risk premium in the case of a constant hazard and constant discounting. For rare events, the hazard rate is likely to be of a smaller order of magnitude than familiar

levels of constant pure rates of time preference. In these cases, the equation shows that the risk premium is likely to be small. The formula for the risk premium is considerably more complicated when both the hazard rate and the pure rate of time preference are non-constant; in these cases the risk premium can be large. Thus, this model provides an example of a situation where hyperbolic discounting is potentially very significant.

In the Ramsey model, the social discount rate equals $\rho + \eta g$, where ρ , η and g are, respectively, the pure rate of time preference, the elasticity of marginal utility, and the growth rate of consumption. If one adopts a very small value of ρ , as in the SR, allowing this parameter to change over time is not going to affect the outcome much. Similarly, if ηg is large, for whatever reason, a changing pure rate of time preference does not matter much.

I think that at much of the disagreement over the optimal level of climate policy has to do with different views about the value of g , and about long run developments in technology more generally. If we are reasonably confident that per capita income will continue to grow and that technological advances will make abatement cheaper, and if we are not concerned that future generations will have the same temptation to delay costly actions that we (possibly) face, and if we are not worried about walking off a climate-related cliff, then the advice to phase in climate policy gradually (the “policy ramp”) has considerable appeal.

6 What next?

If I knew the answer to that question, I would be working on that project rather than writing a “synthesis paper”. Instead, I will take this opportunity to recommend two lines of enquiry that I think are creative and which offer the possibility of new insights about climate policy.

Most of our models use a single composite consumption good, and assume that climate damage reduces the productivity of man-made capital. An alternative recognizes that natural capital and man-made capital provide services that are imperfectly substitutable in production, consumption, or both. Hoel and Sterner (2007) lay the groundwork in analyzing this more complicated model; Sterner and Persson (2008) provide numerical evidence of the importance of the model with different types of capital. Traeger (2008) analyzes the role of the different stocks in determining social discount rates and relates the results to models of sustainability.

Most of our models assume that our objective is to maximize the expectation of the discounted sum of welfare in the current and future periods. Traeger (2009) develops the concept of intertemporal risk aversion, which has a simple explanation. Suppose that an agent is indifferent between two (nonconstant) consumption streams, $X = (x_1, x_2, \dots, x_n)$ and $Y = (y_1, y_2, \dots, y_n)$. Denote Z_{\max} and Z_{\min} , as, respectively, the consumption sequences created by taking, element by element, the larger (respectively, the smaller) of the values of x_i and y_i . This agent is intertemporally risk averse if and only if she prefers X (or the welfare-equivalent Y) compared to obtaining Z_{\max} or Z_{\min} with equal probability. In the former case she experiences some good and some bad periods, and with the lottery she experiences only good or only bad periods. The usual type of dynamic model, which assumes additively separable utility across periods, implies intertemporal risk neutrality. Traeger (2009) shows that intertemporal risk aversion can have a significant effect on the social discount rate; he also discusses the relation between intertemporal risk aversion and ambiguity aversion (arising when the distribution of a random variable is not known). Ha-Duong and Triech (2004) illustrate the importance of a non-additively separable model in setting climate change policy.

7 Conclusion

The paper makes four main points:

- The sensitivity of optimal climate policy to discounting assumptions is model-specific. Except for very simple models, we do not have a good basis for intuition about the circumstances that make discounting (very) important.
- The likelihood of intra-generational conflict arising from optimal climate policy may be exaggerated under two circumstances: (i) if we fail to recognize that we can internalize climate damages by changing the composition of investment between man-made and natural capital, or (ii) if we fail to recognize that the correct business-as-usual baseline involves investment decisions that internalize neither future nor current damages.
- Uncertainty about catastrophic events might cause the present value of the marginal utility of the first “sure unit” of consumption transferred

from today, into the future, to be infinite. This observation does not mean that uncertainty about catastrophic events creates a fundamentally new complication in sensibly applying cost-benefit analysis to the question of climate policy. Neither does it imply that standard discounting is a second order issue, in comparison to the uncertainty about catastrophes.

- One-point spatial perspective means not only that distant objects appear smaller, but also that the rate of decrease of apparent size falls with increased distance: our ability to distinguish between distant objects diminishes with distance. If one accepts Ramsey's analogy between perspective in space and in time, then our "natural" way of viewing the distant future involves hyperbolic rather than constant discounting. Hyperbolic discounting requires that we think about climate policy as an intergenerational game, not a standard intertemporal optimization problem.

References

- Foley, D.: 2008, The economic fundamentals of global warming, *in* J. Harris and N. Goodwin (eds), *Twenty-first century macroeconomics: Responding to the Climate Challenge*, Edward Elgar Publishing.
- Fujii, T. and Karp, L.: 2008, Numerical analysis of non-constant pure rate of time preference: a model of climate policy, *Journal of Environmental Economics and Management*, *forthcoming* **56**, 83–101.
- Ha-Duong, M. and Treich, N.: 2004, Risk aversion, intergenerational equity and climate change, *Environmental and Resource Economics* **28**(2), 195–207.
- Hoel, M. and Sterner, T.: 2007, Discounting and relative prices, *Climatic Change* **84**, 265– 80.
- Horowitz, J. and Lange, A.: 2008, What’s wrong with infinity - a note on weitzman’s dismal theorem, University of Maryland Working Paper.
- Karp, L.: 2005, Global warming and hyperbolic discounting, *Journal of Public Economics* **89**, 261–282.
- Karp, L. and Tsur, Y.: 2008, Time perspective, discounting and climate change policy. Unpublished working paper; <http://are.Berkeley.EDU/karp/>.
- Nordhaus, W. D.: 2007, A review of the Stern Review on the economics of climate change, *Journal of Economic Literature* (3), 686 – 702.
- Norhaus, W.: 2009, An analysis of the dismal theorem, Cowles Foundatiuon Discussion Paper No 1686.
- Ramsey, F. P.: 1928, A mathematical theory of saving, *Economic Journal* **38**, 543–559.
- Rezai, A., Foley, D. and Taylor, L.: forthcoming, Global warming and economic externalities, *Economic Theory* **xx**, xx.
- Stern, N.: 2007, *The Economics of Climate Change*, Cambridge University Press.

Sterner, T. and Persson, M.: 2008, An even sterner review: Introducing relative prices into the discounting debate, *Review of Environmental Economics and Policy* **2**, 61–76.

Traeger, C.: 2008, Sustainability, limited substitutability and non-constant social discount rates. DARE Working Paper.

Traeger, C.: 2009, The social discount rate under intertemporal risk aversion and ambiguity. DARE Working Paper.

Weitzman, M.: 2009, Reactions to the Nordhaus Critique, unpublished manuscript.

Weitzman, M.: forthcoming, On modeling and interpreting the economics of catastrophic climate change, *Review of Economics and Statistics* **xx(x)**, xx.

Appendix: Proof of Proposition 1

Proof. If the two parallel segments of a trapezoid have length A and E , as shown in Figure 5, and the distance between them (the height of the trapezoid) is H , then the length of a segment parallel to them, at height h is

$$f(h) = A - \frac{A - E}{H}h \quad (10)$$

If the height of a trapezoid with parallel sides of length z and $x = z + 2\left(\frac{x-z}{2}\right)$ is y , then the height of the intersection of the principal diagonals is (using the property of similar triangles and Figure 6)

$$\begin{aligned} \frac{y}{z + \frac{x-z}{2}} &= \frac{h}{\frac{x}{2}} \Rightarrow \\ h &= y \frac{x}{x+z} \end{aligned} \quad (11)$$

Normalize by setting the height of the large trapezoid in Figure 5 equal to 1: $H = 1$. Using equation (11) I have

$$a + b = \frac{HA}{A + E} = \frac{A}{A + E}. \quad (12)$$

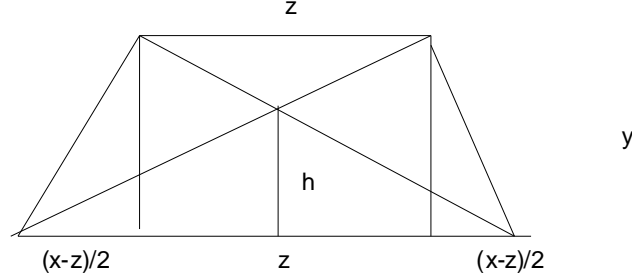


Figure 6:

Using equation (12) in equation (10) I have

$$C = A - \frac{A-E}{H} \frac{A}{A+E} = A - \frac{A-E}{1} \frac{A}{A+E} = 2A \frac{E}{A+E} \quad (13)$$

Using equations (13) and (12) in equation (11) I have

$$a = (a+b) \frac{A}{A+C} = \frac{A}{A+E} \frac{A}{A + (2A \frac{E}{A+E})} = \frac{A}{A+3E}. \quad (14)$$

Equations (12) and (14) imply

$$b = a + b - a = \frac{A}{A+E} - \frac{A}{A+3E} = 2A \frac{E}{A^2 + 4AE + 3E^2} \quad (15)$$

so

$$\frac{b}{a} = \frac{2A \frac{E}{A^2 + 4AE + 3E^2}}{\frac{A}{A+3E}} = 2 \frac{E}{A+E}$$

I now use equation (12) to write

$$c + d = H - (a + b) = 1 - \frac{A}{A+E} = \frac{E}{A+E}. \quad (16)$$

Using equations (11), (13) and (16) to obtain

$$c = \frac{C(c+d)}{C+E} = \frac{2A \frac{E}{A+E} \left(\frac{E}{A+E} \right)}{2A \frac{E}{A+E} + E} = 2A \frac{E}{3A^2 + 4AE + E^2}. \quad (17)$$

Using equations (17) and (15) gives

$$\frac{c}{b} = \frac{2A \frac{E}{3A^2+4AE+E^2}}{2A \frac{E}{A^2+4AE+3E^2}} = \frac{A+3E}{3A+E}$$

which implies

$$\frac{\frac{c}{b} - \frac{b}{a}}{\frac{b}{a}} = \frac{\frac{A+3E}{3A+E} - 2\frac{E}{A+E}}{2\frac{E}{A+E}} = \frac{(A-E)^2}{2E^2+6AE} > 0,$$

thus establishing inequality (i) in Definition 1.

To establish inequality (ii) in the Definition, use equations (10) and (14) to obtain.

$$B = A - (A-E)a = A - (A-E)\frac{A}{A+3E} = 4A\frac{E}{A+3E}. \quad (18)$$

Using equation (13) and (18):

$$\begin{aligned} \frac{B}{A} &= \frac{4A\frac{E}{A+3E}}{A} = 4\frac{E}{A+3E} \\ \frac{C}{B} &= \frac{2A\frac{E}{A+E}}{4A\frac{E}{A+3E}} = \frac{(A+3E)}{2(A+E)} \end{aligned}$$

so

$$\frac{\frac{C}{B} - \frac{B}{A}}{\frac{C}{B}} = \frac{1}{AC} (AC - B^2) = \frac{A2A\frac{E}{A+E} - \left(4A\frac{E}{A+3E}\right)^2}{A2A\frac{E}{A+E}} = \frac{(A-E)^2}{(A+3E)^2} > 0$$

□

CESifo Working Paper Series

for full list see www.cesifo-group.org/wp

(address: Poschingerstr. 5, 81679 Munich, Germany, office@cesifo.de)

- 2697 Steven Brakman, Charles van Marrewijk and Arjen van Witteloostuijn, Market Liberalization in the European Natural Gas Market – the Importance of Capacity Constraints and Efficiency Differences, July 2009
- 2698 Huifang Tian, John Whalley and Yuezhou Cai, Trade Sanctions, Financial Transfers and BRIC's Participation in Global Climate Change Negotiations, July 2009
- 2699 Axel Dreher and Justina A. V. Fischer, Government Decentralization as a Disincentive for Transnational Terror? An Empirical Analysis, July 2009
- 2700 Balázs Égert, Tomasz Koźluk and Douglas Sutherland, Infrastructure and Growth: Empirical Evidence, July 2009
- 2701 Felix Bierbrauer, Optimal Income Taxation and Public Goods Provision in a Large Economy with Aggregate Uncertainty, July 2009
- 2702 Marc Gronwald, Investigating the U.S. Oil-Macroeconomy Nexus using Rolling Impulse Responses, July 2009
- 2703 Ali Bayar and Bram Smeets, Government Deficits in the European Union: An Analysis of Entry and Exit Dynamics, July 2009
- 2704 Stergios Skaperdas, The Costs of Organized Violence: A Review of the Evidence, July 2009
- 2705 António Afonso and Christophe Rault, Spend-and-tax: A Panel Data Investigation for the EU, July 2009
- 2706 Bruno S. Frey, Punishment – and beyond, July 2009
- 2707 Michael Melvin and Mark P. Taylor, The Crisis in the Foreign Exchange Market, July 2009
- 2708 Firouz Gahvari, Friedman Rule in a Model with Endogenous Growth and Cash-in-advance Constraint, July 2009
- 2709 Jon H. Fiva and Gisle James Natvik, Do Re-election Probabilities Influence Public Investment?, July 2009
- 2710 Jarko Fidrmuc and Iikka Korhonen, The Impact of the Global Financial Crisis on Business Cycles in Asian Emerging Economies, July 2009
- 2711 J. Atsu Amegashie, Incomplete Property Rights and Overinvestment, July 2009

- 2712 Frank R. Lichtenberg, Response to Baker and Fugh-Berman's Critique of my Paper, "Why has Longevity Increased more in some States than in others?", July 2009
- 2713 Hans Jarle Kind, Tore Nilssen and Lars Sørgard, Business Models for Media Firms: Does Competition Matter for how they Raise Revenue?, July 2009
- 2714 Beatrix Brügger, Rafael Lalive and Josef Zweimüller, Does Culture Affect Unemployment? Evidence from the Röstigraben, July 2009
- 2715 Oliver Falck, Michael Fritsch and Stephan Heblich, Bohemians, Human Capital, and Regional Economic Growth, July 2009
- 2716 Wladimir Raymond, Pierre Mohnen, Franz Palm and Sybrand Schim van der Loeff, Innovative Sales, R&D and Total Innovation Expenditures: Panel Evidence on their Dynamics, July 2009
- 2717 Ben J. Heijdra and Jochen O. Mierau, Annuity Market Imperfection, Retirement and Economic Growth, July 2009
- 2718 Kai Carstensen, Oliver Hülsewig and Timo Wollmershäuser, Price Dispersion in the Euro Area: The Case of a Symmetric Oil Price Shock, July 2009
- 2719 Katri Kosonen and Gaëtan Nicodème, The Role of Fiscal Instruments in Environmental Policy, July 2009
- 2720 Guglielmo Maria Caporale, Luca Onorante and Paolo Paesani, Inflation and Inflation Uncertainty in the Euro Area, July 2009
- 2721 Thushyanthan Baskaran and Lars P. Feld, Fiscal Decentralization and Economic Growth in OECD Countries: Is there a Relationship?, July 2009
- 2722 Nadia Fiorino and Roberto Ricciuti, Interest Groups and Government Spending in Italy, 1876-1913, July 2009
- 2723 Andreas Wagener, Tax Competition, Relative Performance and Policy Imitation, July 2009
- 2724 Hans Fehr and Fabian Kindermann, Pension Funding and Individual Accounts in Economies with Life-cyclers and Myopes, July 2009
- 2725 Ernesto Reuben and Arno Riedl, Enforcement of Contribution Norms in Public Good Games with Heterogeneous Populations, July 2009
- 2726 Kurt Schmidheiny and Marius Brülhart, On the Equivalence of Location Choice Models: Conditional Logit, Nested Logit and Poisson, July 2009
- 2727 Bruno S. Frey, A Multiplicity of Approaches to Institutional Analysis. Applications to the Government and the Arts, July 2009

- 2728 Giovanni Villani, A Strategic R&D Investment with Flexible Development Time in Real Option Game Analysis, July 2009
- 2729 Luca Di Corato and Michele Moretto, Investing in Biogas: Timing, Technological Choice and the Value of Flexibility from Inputs Mix, July 2009
- 2730 Gilad D. Aharonovitz, Nathan Skuza and Faysal Fahs, Can Integrity Replace Institutions? Theory and Evidence, July 2009
- 2731 Michele Moretto and Sergio Vergalli, Managing Migration through Conflicting Policies: an Option-theory Perspective, July 2009
- 2732 Volker Nitsch, Fly or Cry: Is Airport Noise Costly?, July 2009
- 2733 Francesco Cinnirella and Joachim Winter, Size Matters! Body Height and Labor Market Discrimination: A Cross-European Analysis, July 2009
- 2734 Samuel Bowles and Sandra Polanía Reyes, Economic Incentives and Social Preferences: A Preference-based Lucas Critique of Public Policy, July 2009
- 2735 Gary Burtless, Lessons of the Financial Crisis for the Design of National Pension Systems, July 2009
- 2736 Helmuth Cremer, Firouz Gahvari and Pierre Pestieau, Fertility, Human Capital Accumulation, and the Pension System, July 2009
- 2737 Hans Jarle Kind and Frank Stähler, Market Shares in Two-Sided Media Industries, July 2009
- 2738 Pamela Campa, Alessandra Casarico and Paola Profeta, Gender Culture and Gender Gap in Employment, August 2009
- 2739 Sebastian Gechert, Supplementary Private Health Insurance in Selected Countries: Lessons for EU Governments?, August 2009
- 2740 Leif Danziger, Endogenous Monopsony and the Perverse Effect of the Minimum Wage in Small Firms, August 2009
- 2741 Yan Dong and John Whalley, A Third Benefit of Joint Non-OPEC Carbon Taxes: Transferring OPEC Monopoly Rent, August 2009
- 2742 Valentina Bosetti, Carlo Carraro and Massimo Tavoni, Climate Change Mitigation Strategies in Fast-Growing Countries: The Benefits of Early Action, August 2009
- 2743 Christina Felfe, The Willingness to Pay for Job Amenities: Evidence from Mothers' Return to Work, August 2009
- 2744 Jörg Franke, Christian Kanzow, Wolfgang Leininger and Alexandra Väh, Effort Maximization in Asymmetric N-Person Contest Games, August 2009

- 2745 Bruno S. Frey and Paolo Pamini, Making World Heritage Truly Global: The Culture Certificate Scheme, August 2009
- 2746 Frank N. Caliendo, Is Social Security behind the Collapse of Personal Saving?, August 2009
- 2747 Caterina Liesegang and Marco Runkel, Corporate Income Taxation of Multinationals and Fiscal Equalization, August 2009
- 2748 Chrysovalantou Milliou and Apostolis Pavlou, Upstream Horizontal Mergers and Efficiency Gains, August 2009
- 2749 Rüdiger Pethig and Christian Wittlich, Interaction of Carbon Reduction and Green Energy Promotion in a Small Fossil-Fuel Importing Economy, August 2009
- 2750 Kai Carstensen, Oliver Hülsewig and Timo Wollmershäuser, Monetary Policy Transmission and House Prices: European Cross-country Evidence, August 2009
- 2751 Olaf Posch, Explaining Output Volatility: The Case of Taxation, August 2009
- 2752 Beatrice Scheubel, Daniel Schunk and Joachim Winter, Don't Raise the Retirement Age! An Experiment on Opposition to Pension Reforms and East-West Differences in Germany, August 2009
- 2753 Daniel G. Arce, Dan Kovenock and Brian Roberson, Suicide Terrorism and the Weakest Link, August 2009
- 2754 Mario Larch and Wolfgang Lechthaler, Comparative Advantage and Skill-Specific Unemployment, August 2009
- 2755 Horst Raff and Nicolas Schmitt, Buyer Power in International Markets, August 2009
- 2756 Seppo Kari, Hanna Karikallio and Jukka Pirttilä, The Impact of Dividend Taxation on Dividends and Investment: New Evidence Based on a Natural Experiment, August 2009
- 2757 Mirco Tonin and Michael Vlassopoulos, Disentangling the Sources of Pro-social Behavior in the Workplace: A Field Experiment, August 2009
- 2758 Nicole Grunewald and Inmaculada Martínez-Zarzoso, Driving Factors of Carbon Dioxide Emissions and the Impact from Kyoto Protocol, August 2009
- 2759 Yu-Fu Chen and Michael Funke, Booms, Recessions and Financial Turmoil: A Fresh Look at Investment Decisions under Cyclical Uncertainty, August 2009
- 2760 Jan-Egbert Sturm and Jakob de Haan, Does Central Bank Communication really Lead to better Forecasts of Policy Decisions? New Evidence Based on a Taylor Rule Model for the ECB, August 2009
- 2761 Larry Karp, Sacrifice, Discounting and Climate Policy: Five Questions, August 2009