PRELIMINARY
DRAFT

(Not to be Quoted)

COST-BENEFIT ANALYSIS AND CLIMATE CHANGE

WP 12.7

Stephen C. Peck, Thomas J. Teisberg

January 1993

Energy Modeling Forum
Terman Engineering Center
Stanford University
Stanford, California

COST BENEFIT ANALYSIS AND CLIMATE CHANGE[1]

Stephen C. Peck
Electric Power Research Institute

and

Thomas J. Teisberg
Teisberg Associates

I. INTRODUCTION

The accumulation of greenhouse gases in the atmosphere raises the possibility of rising global mean surface temperature over the next couple of centuries. Proposed policy responses to global warming have generally been framed in terms of limits on emissions of greenhouse gases, especially CO_2. Most such proposals seem to be quite arbitrary in character, e.g. limit CO_2 emissions to 80 percent of 1990 levels. The very arbitrariness of such proposals raises the question: **what is the best path of emissions over time?**

It is not possible to determine an optimal emissions path without considering both the costs and the benefits associated with reductions in CO_2 emissions. The costs of reducing emissions have received a considerable amount of attention from policy analysts; many of their analyses are represented in this Energy Modeling Forum study. However, much less attention has been paid to estimating the benefits of emission reduction (or equivalently, the avoided warming cost).

Nordhaus (1990) has attempted to synthesize available information regarding the cost to the US of a three degree Celsius increase in mean surface temperature. Based on this work, Nordhaus concludes that measurable costs of warming are in the neighborhood of one-quarter of one percent of GDP. However, there are important possible costs of warming that are excluded from this total, for lack of any information regarding their magnitude. These include damage to unmanaged natural systems including loss of ecosystems and species, and losses in the amenity values of everyday life. Nordhaus offers a guess that including these omitted costs and considering the world as a whole, the total cost could be one or two percent of global income.

While Nordhaus's estimate is a useful starting point, important questions remain: What are the costs associated with a five, six, or eight degree rise in temperature? Does it matter whether temperature rise occurs overnight or over two centuries? These are exceedingly difficult questions to answer, but they need to be answered if CO_2 emission limits are to be based on a balance between the costs and benefits of emission reductions.

Obviously, we cannot begin to resolve these questions here. What we can do, however, is offer some insight into the possible character of emission limits or carbon taxes, if these limits or taxes are determined from a balancing of costs and benefits; in addition, we can offer some insight into what the benefits of emission reduction might have to be in order for the kinds of policies now being discussed to be at least roughly correct, taking account of both costs and benefits.

ANALYTICAL FRAMEWORK

To explore the implications of alternative damage functions for optimal emissions and carbon tax paths, we use the CETA model. CETA represents world-wide economic growth, energy consumption, energy technology choice, global warming, and global warming costs.[2] Figure 1 presents a schematic overview of the key relationships in the model. Energy technologies and the oil, gas, and coal resource bases are inputs to an energy submodel, which supplies energy inputs to a production submodel, and the CO_2 by-product to the warming submodel. In the production submodel, energy, labor, and capital inputs are used to produce output which is then allocated to consumption, investment, energy costs, and damage costs of warming. Because energy consumption and energy technology choices are considered together

Figure 1: CETA Model Overview

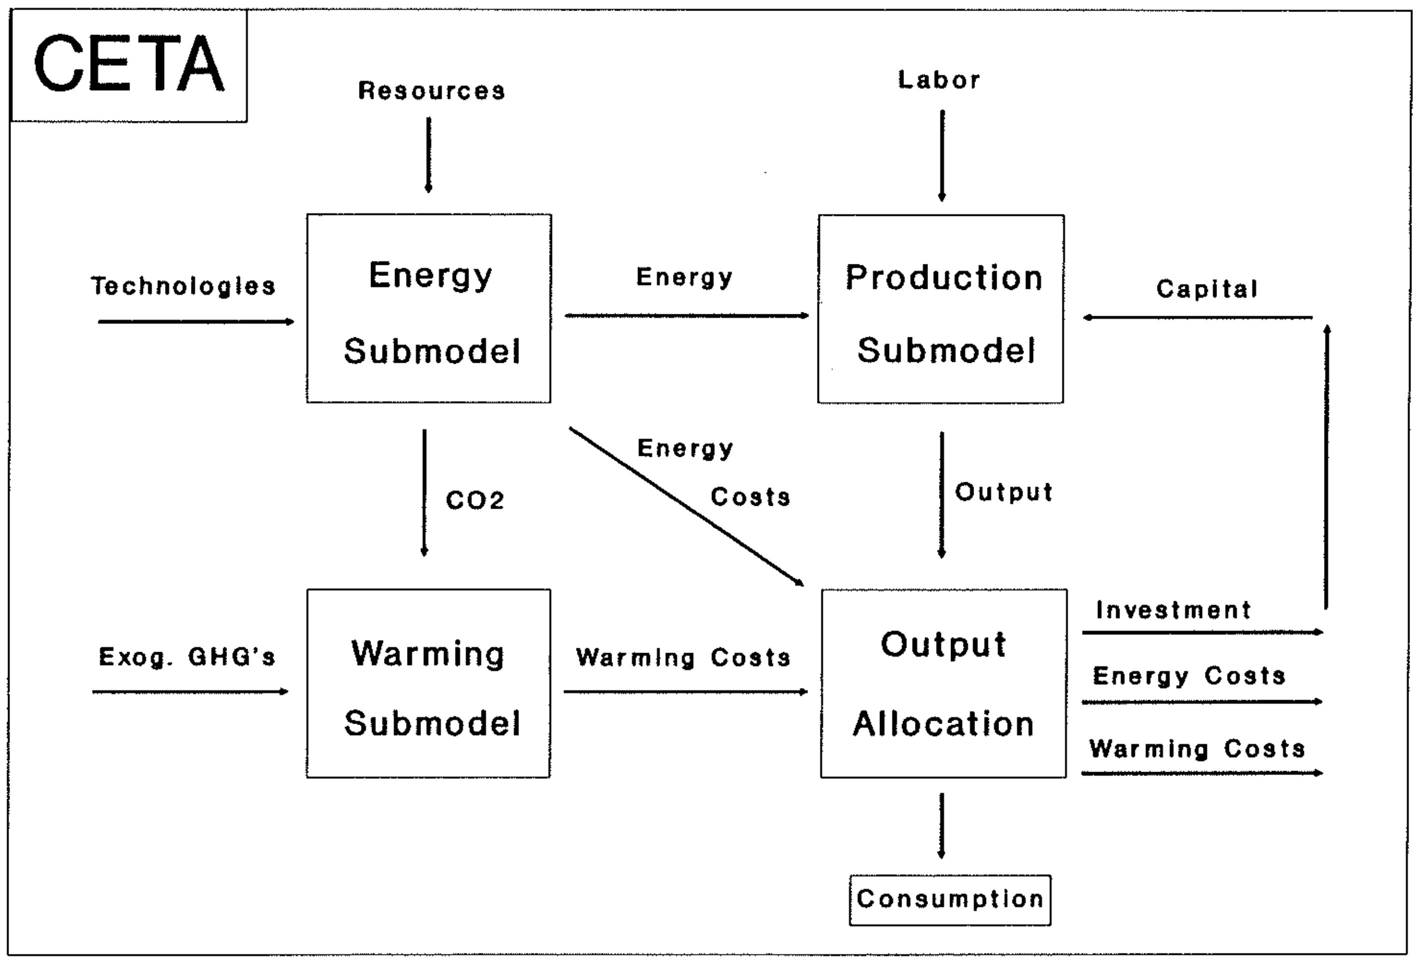

with warming damage costs, the time paths of CO_2 emissions and carbon taxes in our model reflect an optimal balancing of the cost of emission reduction and the benefit of reduced global warming.

In our work with CETA, we have looked at the implications of damage functions that are defined both on the level of temperature and on its rate of change. Also, for each type of damage function, we have considered the implications of damage which is a linear or a power function of temperature or temperature change.

The key difference between the level-related and rate-related damage functions is in the timing of the costs of warming. With the level-related function, damages start low, and rise over time as temperature change increases; moreover, when temperature change is ultimately stabilized, these costs persist. With the rate-related function, damages tend to be higher initially, since the rate of temperature change is high before the actual change becomes high. However, as temperature is ultimately stabilized at a higher level, the rate-related damage cost goes down, ultimately to zero.

For the case where damage depends on the level of temperature rise, the damage functions we have considered are of the form:

(1) $$D_t = \alpha \cdot L_t \cdot T_t^{\lambda}$$

where
D_t = *annual damage, time t*
α = *a scaling constant*
L_t = *labor index* (L_1 = 1.0)
T_t = *temperature level (above pre-industrial)*
λ = *power of the damage function.*

Scaling the damage function by the labor index, Lt, causes damages at a given temperature rise to remain roughly the same proportion of gross production, as production rises over time.

When damage is a function of the rate of temperature change, we use the following function defined on past decadal changes in temperature:

(2) $$D_t = \alpha(L_t \cdot [\Delta T]^{\lambda}_{-0} + (1-\mu) \cdot L_{t-1} \cdot [\Delta T]^{\lambda}_{-1} + (1-\mu)^2 \cdot L_{t-2} \cdot [\Delta T]^{\lambda}_{-2} + \ ... \ )$$

where D_t = *annual damage in the current decade*
α = *scaling constant*
L_t = *labor index* (L_1=1.0)
$[\Delta T]_{-t}$ = *temperature change, t decades earlier*
λ = *power of the damage function*
μ = *damage recovery rate.*

For modeling convenience, we use the rate-related damage function in the equivalent alternative form:

(3) $$D_t = \alpha \cdot L_t \cdot [\Delta T]^{\lambda}_{-0} + (1-\mu) \cdot D_{-1}$$

where D_{-1} = *annual damage in the previous decade.*

The parameter, λ, in both types of damage functions determines the power of the function. When $\lambda=1$, the function is linear, and when $\lambda>1$, the function is non-linear increasing. For results presented here, we use two alternative values for λ: 1 (linear) and 3 (cubic).

In the case of the level-related damage function, the scaling constant, α, is set to calibrate the damage function to some benchmark estimate of warming costs. For most of the results presented here, α is set so that damages in 1990 would be two percent of 1990 gross world production, at a three degree temperature rise. This benchmark corresponds to the highest estimate of damage in Nordhaus (1990). In addition, we show some results here of experiments in which we scale up α so that emissions paths and carbon taxes are roughly comparable to those for a policy of stabilizing emissions at the 1990 level.

For the rate-related damage function, it is more difficult to choose a scaling constant, since there is no available damage estimate premised on any particular rate of warming. Yet, as temperature rises three degrees over the next century or so, the projected rate of increase in temperature is approximately 0.25 degrees per decade;[3] thus one possibility would be to associate the our benchmark damage for a three degree warming with this rate of increase. However, we think it more reasonable to associate this level of damages with a slower rate of increase (or equivalently, to associate higher damages with the 0.25 degrees per decade rate of increase). Thus we somewhat arbitrarily choose to set the scaling constant α so that damages are two percent of 1990 gross world production, for a constant temperature rise of 0.20 degrees per decade (and a constant labor index, L_t).[4]

In the rate-related damage function, the parameter, μ, represents the potential for past damage to be reduced over time as ecosystems adjust to higher temperature. In fact, results obtained with this function are not particularly sensitive to the value of μ.[5] For the results presented here, we assume that $\mu=0.1$; this is consistent with a damage recovery of 65 percent after about 220 years.[6]

LINEAR DAMAGE FUNCTIONS

The simplest assumption to make regarding damages is that they are proportional to temperature change or to the rate of temperature change. This assumption means that the power, λ, in the damage function discussed above is 1.

Figure 2 compares carbon emissions in the uncontrolled case with optimal emissions in the two cases where level-related and rate-related linear damage functions are assumed. It is striking how little difference there is between any of these emission paths. Apparently, with linear damages it makes almost no difference whether damages are assumed to depend on the

Figure 2: Carbon Emissions

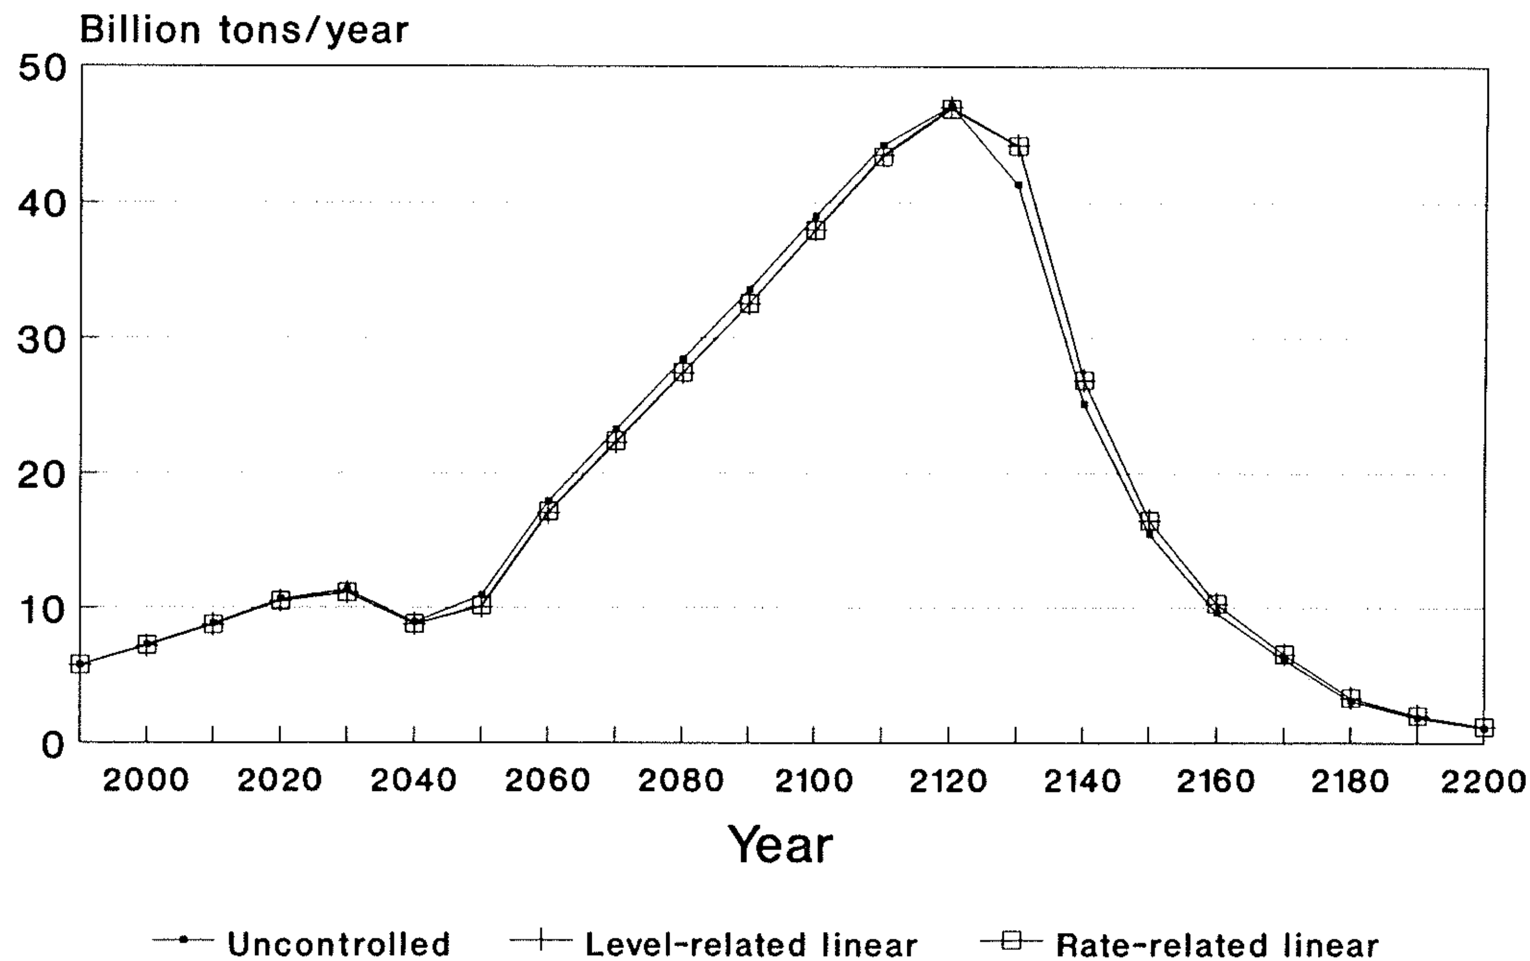

level of temperature or its rate of increase; and in either case, the optimal emission reduction is small, and far short of the emission reduction implied by a policy of stabilizing emissions at the 1990 level or reducing them to some percentage below 1990 levels.

CUBIC DAMAGE FUNCTIONS

Rather than being proportional to temperature change or to the rate of temperature change, damages may rise at an accelerating rate as temperature or its rate of change becomes higher. We represent this case with the cubic damage functions, in which the power, λ, of the damage functions is 3. A power of 3 means that damage roughly doubles as a percent of gross production, when the temperature change or its (steady) rate of change goes up by 25 percent.

Figure 3 compares carbon emissions in the uncontrolled case with optimal emissions in the two cases where damages are a cubic function of temperature change or its rate of change. Here there are substantial differences between the uncontrolled emissions path and the two optimal paths. Evidently, damages from warming must be non-linearly increasing in order for significant CO_2 emission reduction to be optimal. Yet, even for these damage functions, the optimal emissions paths for the next century at least remain much above the emission paths that are implied by policies calling for stabilization of emissions at or below the 1990 level.

Figure 3 also reveals a difference in the emission paths for the level-related damage function and that for the rate-related damage function. Viewed over the entire 210 year period shown, the rate-related optimal emissions path is more nearly constant over time. To understand why this is so, consider Figure 4, which shows the time paths of temperature resulting from uncontrolled emissions and those resulting from the optimal emissions paths for each damage function. Since temperature continues to rise over the entire time period, damages are always increasing with the level-related function, and it is optimal to reduce emissions more and more

Figure 3: Carbon Emissions

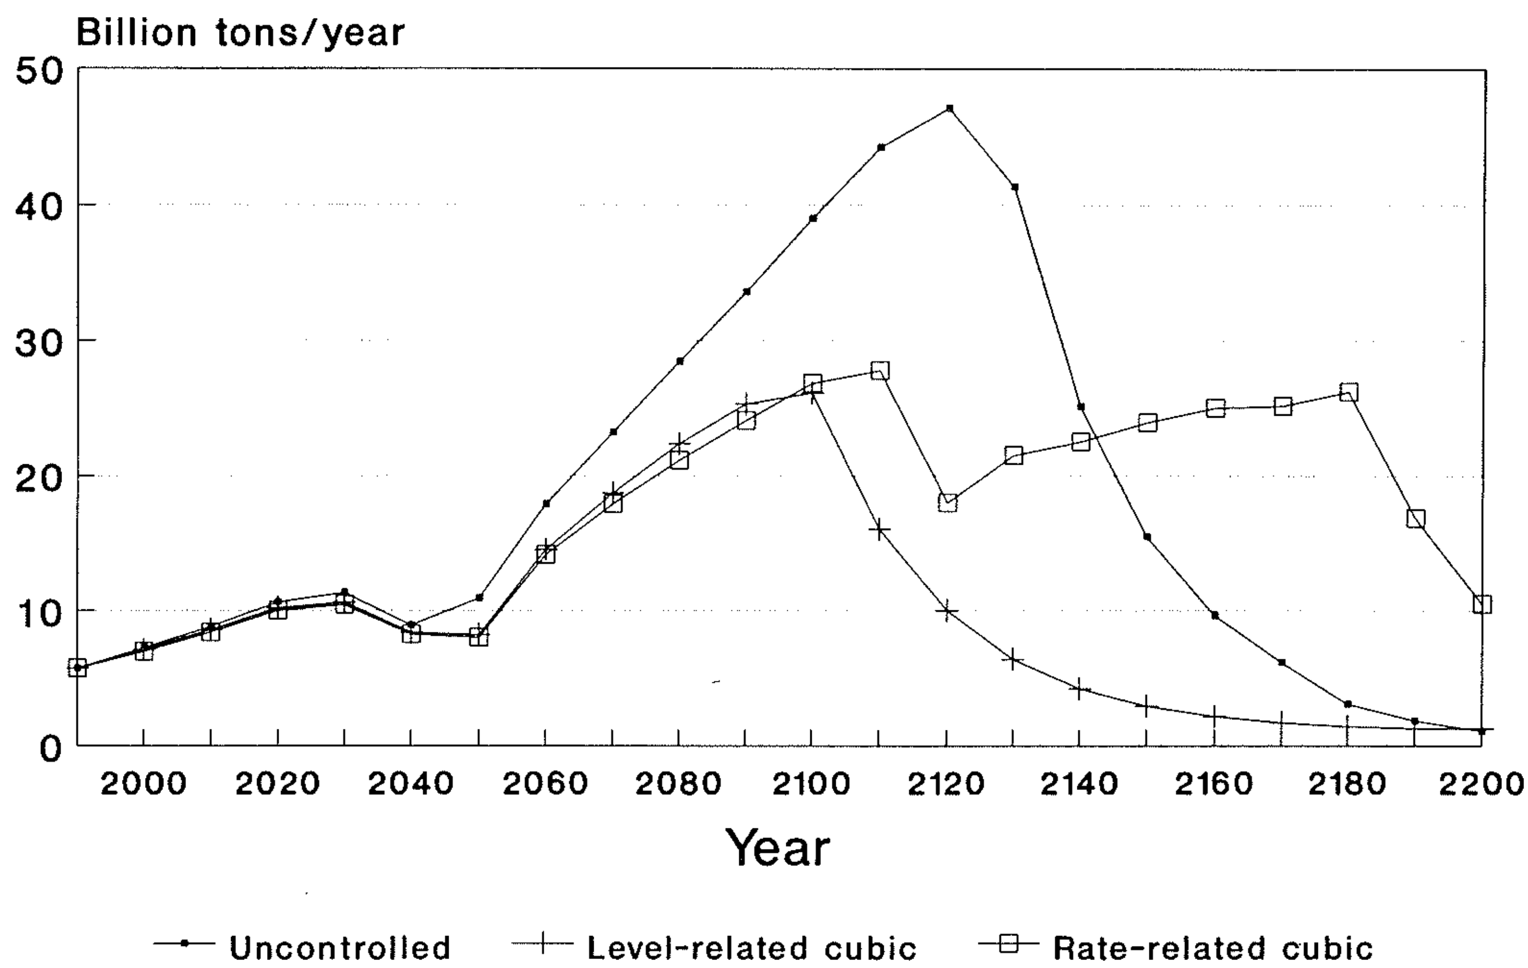

Figure 4: Temperature Rise

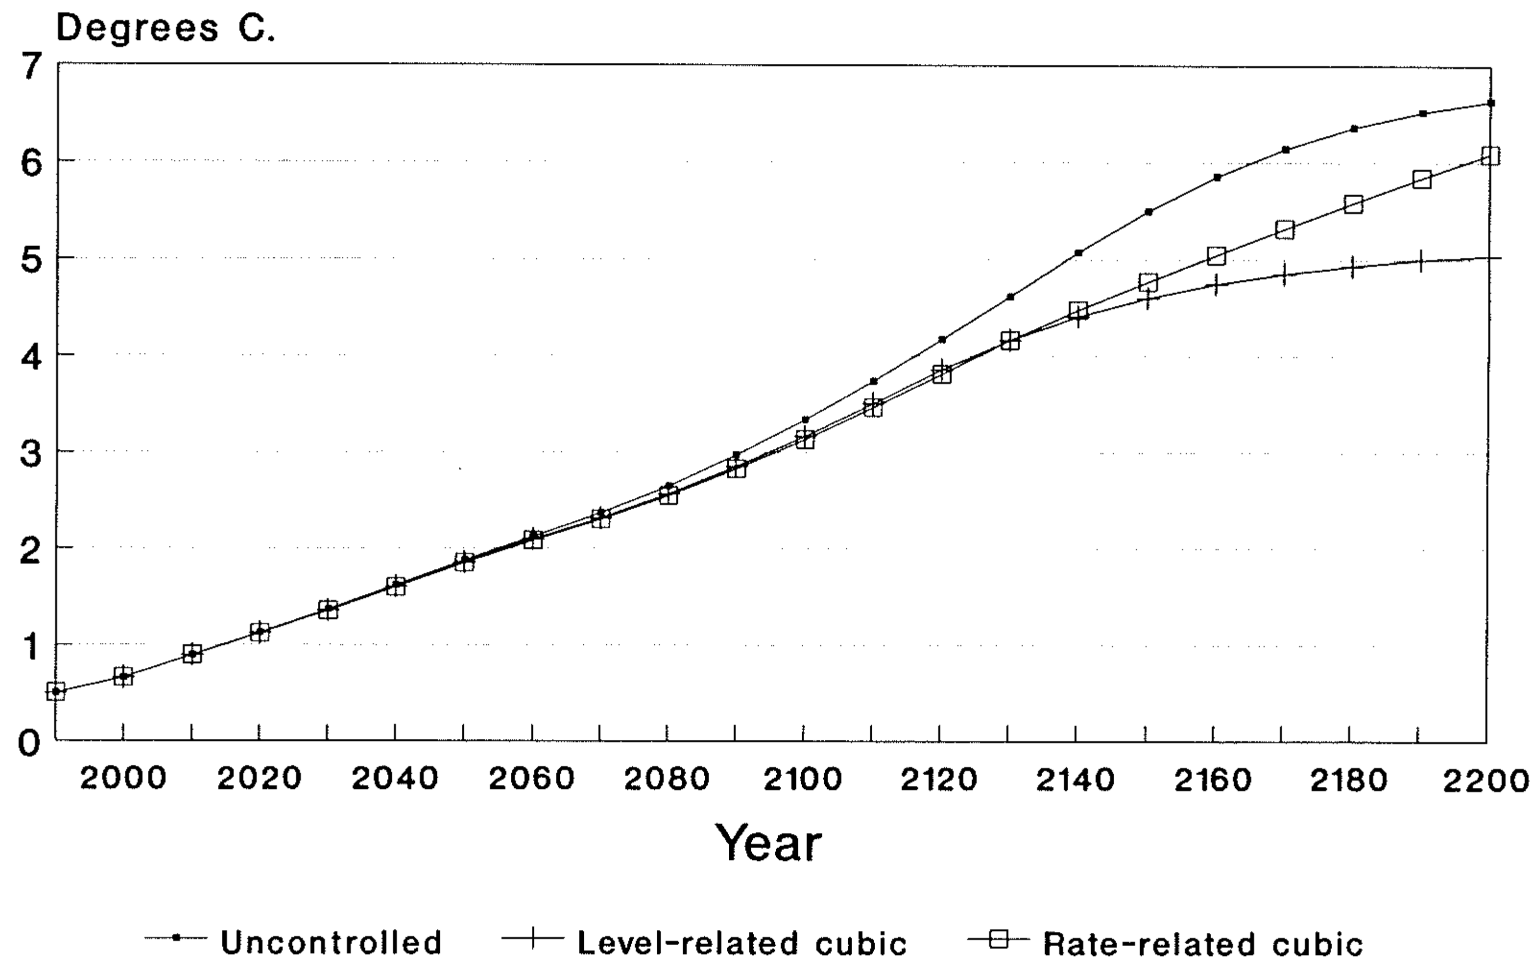

as time goes on. On the other hand, the rate of temperature change along the rate-related damage function temperature path is roughly constant, which means that marginal cost for the rate-related damage function is more nearly constant over time, resulting in more nearly constant optimal emissions.

PROPOSED EMISSION LIMIT POLICIES

Stabilization of emissions at or below 1990 levels is a commonly discussed response to the global warming problem. Indeed, most of the emission reduction scenarios used in the EMF-12 study are intended to represent this type of policy. This leads us to inquire into what damage assumptions would be consistent, or at least roughly consistent, with such policies.

Based on our observations in the preceding sections, linear damage functions could be consistent with these policies only if the scaling constants α were set at extremely high levels relative to the Nordhaus benchmark of two percent of GDP for a three degree warming. Thus we restrict our search for roughly consistent damage functions to the cubic functions. In addition, based on the preceding results, we should expect that rate-related damage functions in general would provide a better fit to policies calling for constant emissions paths over time. To demonstrate this, we present results for a range of damage function types: a pure level-related function, a pure rate-related function, and a mixed level-related and rate-related function.

We will take stabilization of emissions at 6 billion tons per year (roughly the 1990 level) as the policy we approximate with optimal policies for level-related, rate-related, and mixed damage functions. The carbon tax path that would be required to implement a 6 billion tons per year policy is shown in Figure 5. We note that this path displays odd up-and-down cycles in the early years, reflecting increasing and decreasing marginal cost of meeting the 6 billion tons per year emission limit. Starting in 2080, however, this path becomes level at about \$208 per ton;

Figure 5: Carbon Tax

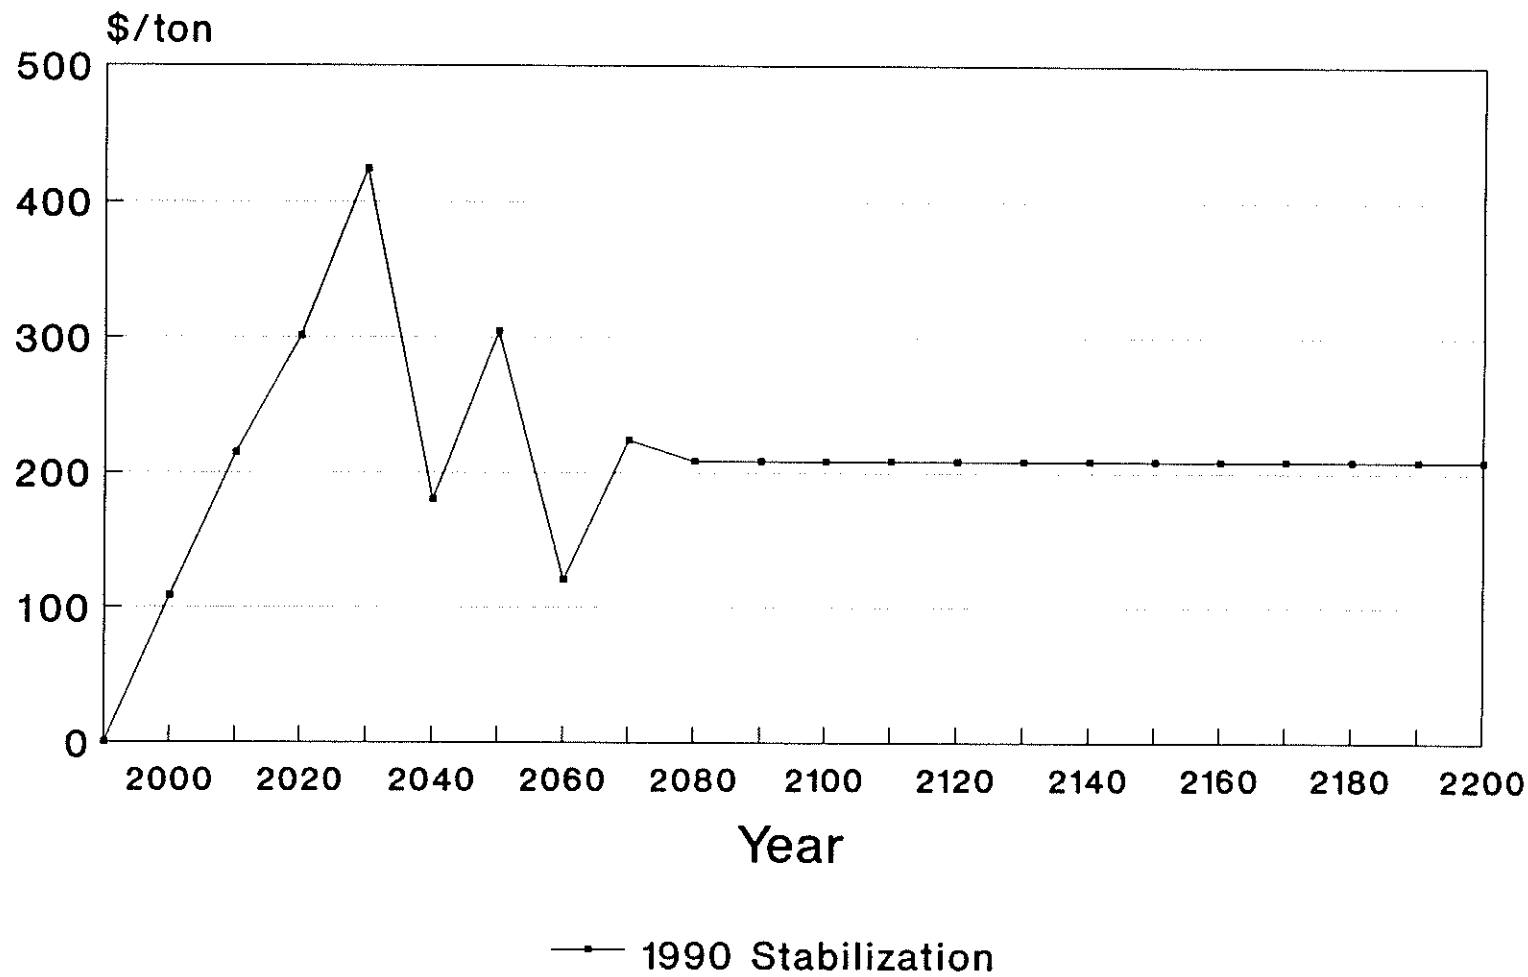

this is the tax that is necessary to equalize the cost of (CO_2 intensive) synfuels and the (carbon-free) non-electric backstop technology.

To find roughly consistent level-related, rate-related, and mixed cubic damage functions, we increase the scaling constant (α) for each type of function until optimal emissions for each average 6 billion tons per year between 1990 and 2100. For the level-related damage function, this requires a scaling constant approximately 3.2 times that for the Nordhaus benchmark; thus damages at three degrees Celsius would have to be about 6.4 percent, rather than 2 percent, of gross world production. For the rate-related damage function, the scaling constant must be raised by a factor of 5.4, implying that damages at a constant 0.20 rate of temperature change (maintained for several decades) would be 10.8 percent of gross world production. Finally, for the mixed level-related and rate-related function, the scaling constants each must be increased by a factor of 2, implying that damages would be approximately 8 percent of gross production at a temperature rise of three degrees and a temperature rate of change of 0.20 degrees per decade.

Figure 6 compares the paths of carbon emissions for the 6 billion tons per year limit policy and for the three roughly consistent damage functions. As the figure shows, optimal emissions for the three damage functions are variable over time, in contrast to the smooth path implied by the stabilization policy. Also, as we had expected, emissions for the rate-related damage function are a better fit to those under the stabilization policy, since they are more nearly constant over time. By contrast, with the level-related damage function, optimal emissions tend to be higher earlier and lower later, relative to the stabilization policy. Emissions for the mixed rate-related and level-related damage function are intermediate between those for the pure rate-related and pure level-related functions.

Figure 7 compares the carbon tax paths for the stabilization policy and for the three roughly consistent damage functions. Here again, carbon tax path for the rate-related damage

Figure 6: Carbon Emissions

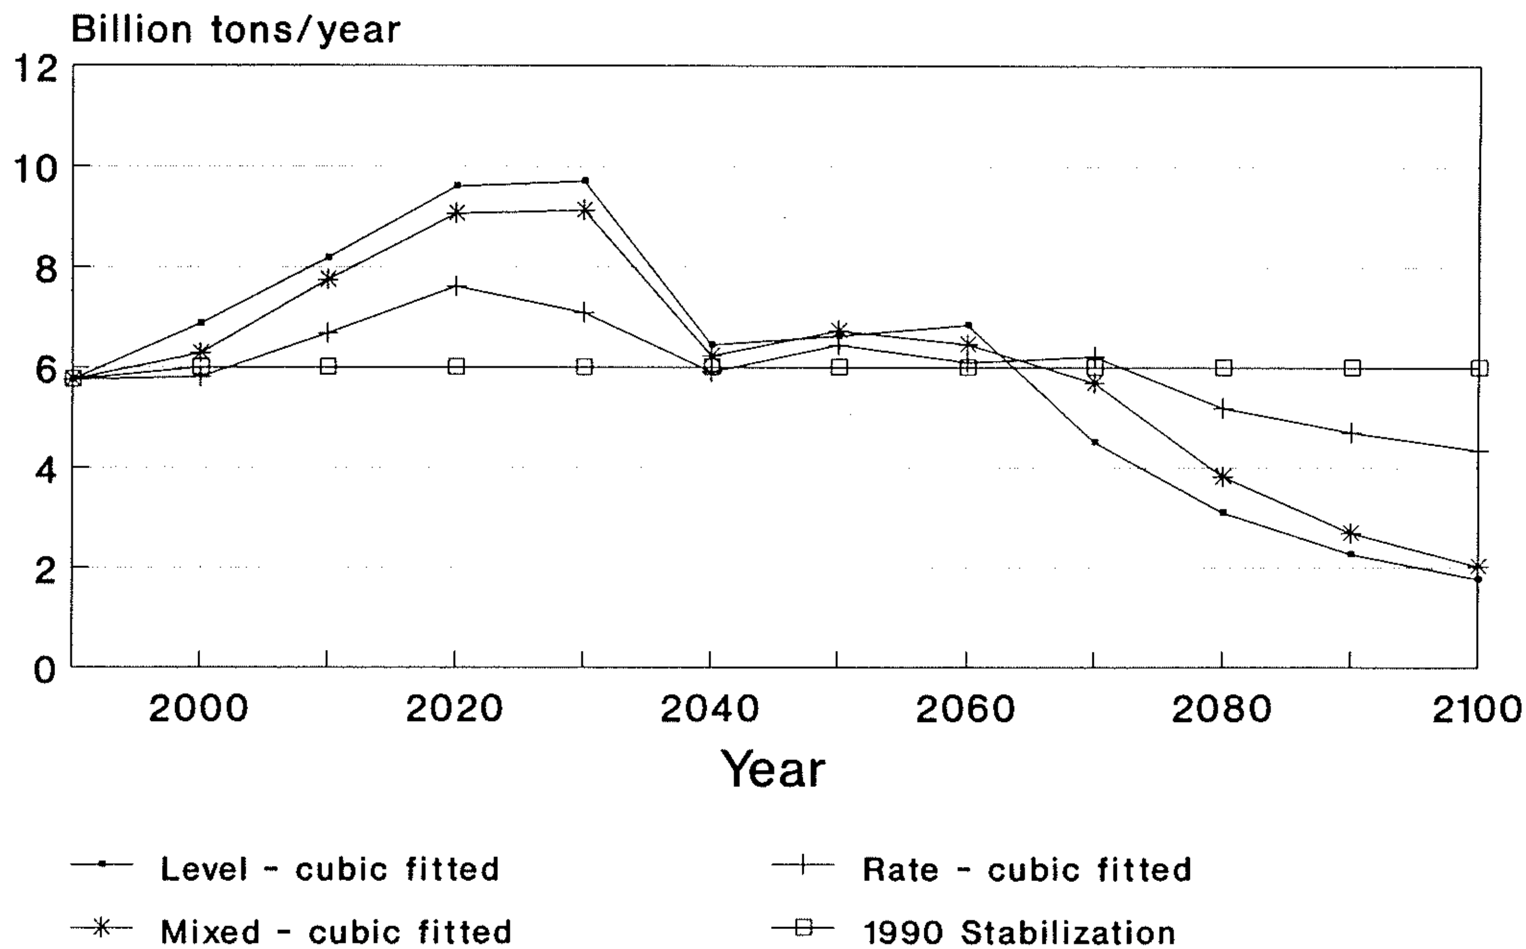

Figure 7: Carbon Tax

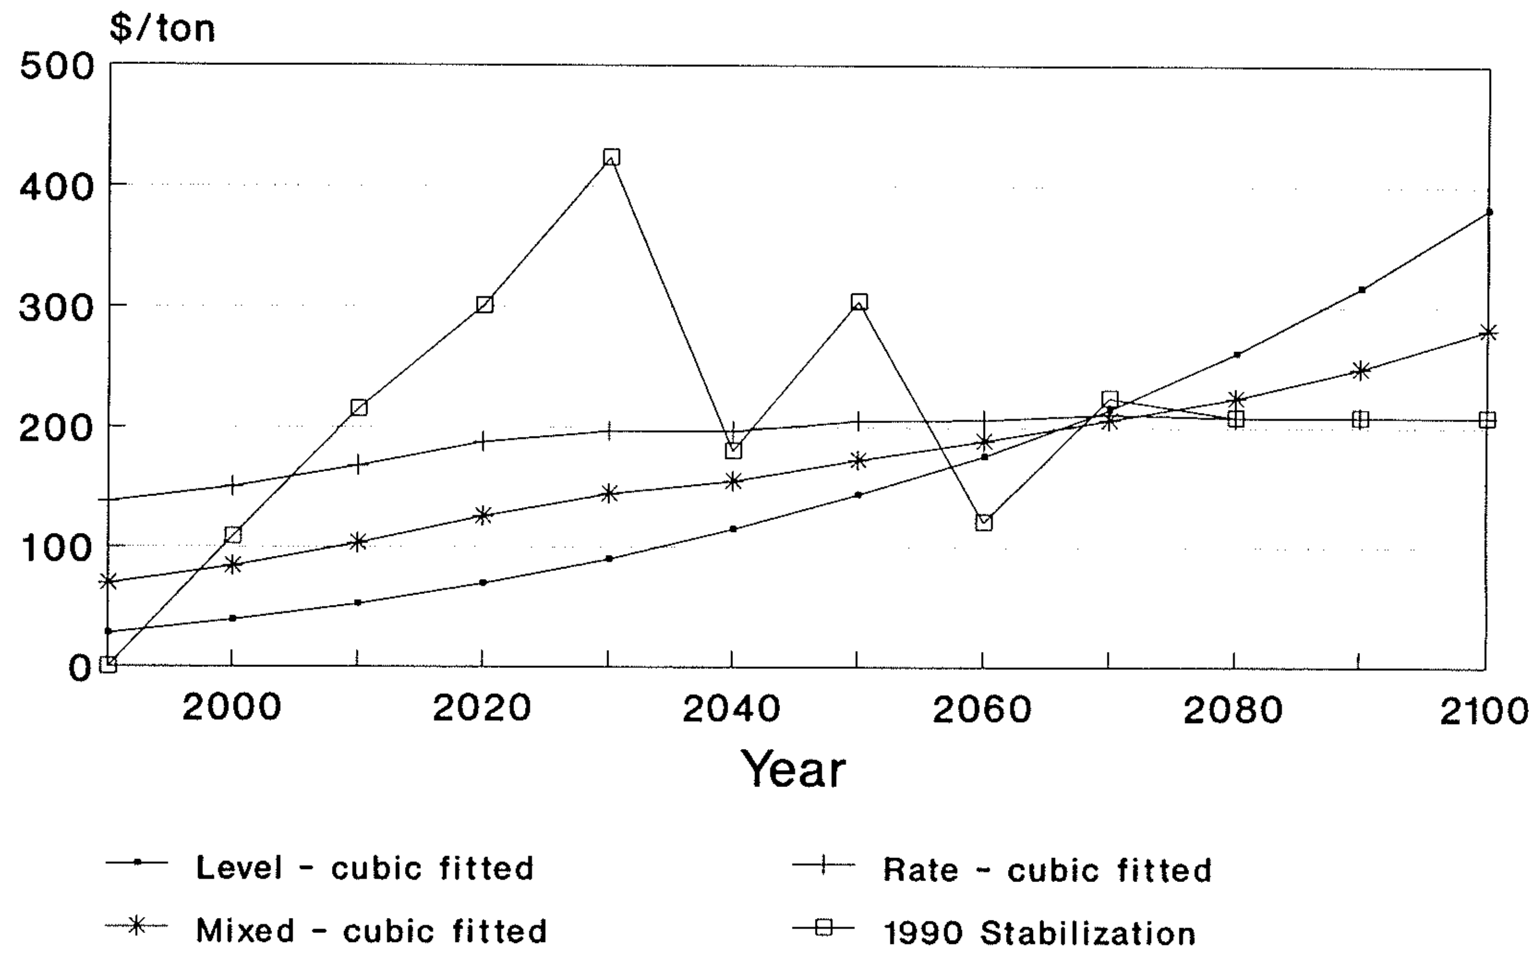

function is the best fit to the carbon tax associated with the stabilization policy. The carbon tax for the level-related damage function does not fit well since it is low initially but then rises continually as temperature rises.[7] Not surprisingly, the carbon tax for the mixed rate-related and level-related damage function is intermediate between those for the pure level-related and the pure rate-related functions. Finally, we note that the optimal carbon tax paths for all of the damage functions are fairly smooth relative to the up-and-down cycles of the carbon tax associated with the stabilization policy.

CONCLUSION

In order for CO_2 emission reduction policy to be based on costs and benefits of reduction, much more needs to be learned about the costs that would result from higher temperatures. However, even before such information becomes available, a model like CETA can be used to explore the implications of alternative possible damage functions and to determine what kinds of damage assumptions would be required to justify the kinds of emission reduction policies now being discussed in the international community.

The model results presented here indicate that non-linearity of the damage function is of central importance, whether damage is a function of temperature change or its rate of change; indeed, if damage is linear, little emission reduction is called for, whether damages are related to temperature change or to its rate of change. On the other hand, if damage is a cubic function of temperature or its rate of change, then substantial emission reductions are appropriate. However, these emission reductions remain well short of those implied by calls to stabilize emissions at or below the 1990 level. In order for the reductions implied by CO_2 emission stabilization to be roughly optimal, level-related or rate-related damages must be cubic and approximately three to five times higher than our benchmarks of 2 percent of gross production for a 3 degree warming and 2 percent of gross production for 0.20 temperature rise per decade warming rate.

Finally, there are interesting differences between the carbon tax and emission paths implied by the stabilization policy and those implied by the optimal policies for the roughly consistent damage functions. Optimal policies generally involve emissions paths with an irregular pattern over time, and carbon taxes that rise smoothly over time. The stabilization policy, on the other hand, implies a smooth path of emissions and an irregular carbon tax.

Optimal carbon taxes that rise smoothly reflect the fact that the marginal benefits of temperature reduction must change smoothly over time, rather than cycling up and down. This is true because a ton of CO_2 emitted this year raises temperature over a span of one or two centuries, while a ton emitted next year raises temperature over a similar span, shifted one year forward. Since these spans are substantially the same, the marginal damage from warming must be substantially the same from one year to the next.

Optimal emissions that vary over time are also sensible. As resources like oil and gas are exhausted, and as new technologies become available, the marginal cost of emission reduction changes. Thus the optimal pattern of reductions reflects greater reductions when such reductions are cheaper, and smaller reductions when they are more expensive.

If we suppose that the true damage function is in fact one of the three roughly consistent damage functions, then there would be an economic loss that would result from use of the stabilization policy, rather than the optimal policy, even though both policies produce the same average reduction in emissions over the time period from 1990 to 2100. If the true damage function were the roughly consistent level-related one, this loss would be about 3.3 trillion dollars; if the true damage function were the mixed rate-related and level-related function, the loss would be lower, about 1.8 trillion dollars; finally, if the true damage function were the best-fitting rate-related function, the loss would be smallest at about 0.7 trillion dollars.[8] In large part, these losses represent the costs of a policy which does not allow emission reductions to cycle up and down as the costs of reduction change over time.[9]

REFERENCES

Maier-Reimer, E. and K. Hasselmann (1987), "Transport and Storage of CO_2 in the Ocean -- An Inorganic Ocean-Circulation Carbon Cycle Model," Climate Dynamics 2, pp. 63-90.

Nordhaus, William D. (1982), "How Fast Should We Graze the Global Commons?," American Economic Review, May, pp. 242-246.

Nordhaus, William D. (1990), "To Slow or Not to Slow: The Economics of the Greenhouse Effect," February 5.

Peck, Stephen C. and Thomas J. Teisberg (1992), "CETA: A Model for Carbon Emissions Trajectory Assessment," The Energy Journal, Vol. 13, No. 1, pp. 55-77.

Peck, Stephen C. and Thomas J. Teisberg (1993), "Temperature Change Related Damage Functions: A Further Analysis with CETA," Energy and Resources (forthcoming).

ENDNOTES

1. The authors are grateful to alan Manne and Rich Richels for advice and help, and to Don Rosenthal for suggesting the focus of this paper. This paper does not represent the views of EPRI or of its members.

2. See Peck and Teisberg (1992) for a detailed description of CETA. The results in this paper are obtained using an updated version of CETA. In the new version, the utility function is population times the logarithm of consumption per capita, instead of simply the logarithm of consumption. In addition, we now use a carbon cycle model adopted from Maier-Reimer and Hasselmann (1987), in which carbon emissions fall into five classes with differing atmospheric lifetimes of 2 years to infinity; before we had used a single class with a lifetime of 233 years. These changes have only a small effect on the results reported in this paper.

3. Figure 4, which is discussed later in the text, shows the path of temperature when emissions are uncontrolled.

4. Optimal emissions paths and carbon taxes are sensitive to this assumption; calibrating two percent of gross production to a lower rate of increase significantly lowers optimal emissions and raises optimal carbon taxes. See Peck and Teisberg (1993).

5. See Peck and Teisberg (1993).

6. We note that if μ (but not α) is changed, damages would no longer equal two percent of gross world production in the steady-state. The following alternative specification of the rate-related function would make steady-state damage invariant to μ:

$$D_t = \alpha \cdot \mu (L_t \cdot [\Delta T]^{\lambda}_{-0} + (1-\mu) \cdot L_{t-1} \cdot [\Delta T]^{\lambda}_{-1} + (1-\mu)^2 \cdot L_{T-2} \cdot [\Delta T]^{\lambda}_{-2} + \ldots)$$

We adopt the specification in the text, since it is more consistent with the interpretation of μ as a damage recovery rate.

7. In previous papers, we have not shown levels of the carbon tax which are in excess of that required to induce substitution of the non-electric backstop for coal based synthetic fuels (i.e. $208 per ton). Here, however, we want to emphasize the implications of the level related cubic damage function for the carbon tax, as temperature continues to rise over time.

8. Of course, since these roughly consistent damage functions were specifically chosen to have optimal emissions similar to the emissions implied by a stabilization policy, the losses from use of the stabilization policy are comparatively small. If the true damage function were something quite different from the roughly consistent functions, the loss from the stabilization policy could be many times greater.

9. These losses also reflect the differences between emissions under the stabilization policy and those under the optimal policies that occur after the year 2100. Since these differences exist far in the future, their present value contribution to the economic losses cited in the text may not be large.