What is the "Damages Function" for Global Warming – and What Difference Might It Make?

Martin L. Weitzman*

July 30, 2009. For *Climate Change Economics.*

Abstract

The existing literature on climate change offers little guidance about why one specification or another of a "damages function" has been selected. Ideally, one wants a functional form that captures reality adequately, yet analytically is sufficiently tractable to yield useful results. This paper gives two plausible risk aversion axioms that a reduced form utility function of temperature change and the capacity to produce consumption might reasonably be required to satisfy. These axioms indicate that the standard practice multiplicative specification of disutility damages from global warming, as well as its additive analogue, are special cases of this paper's theoretically-derived utility function. Empirically, the paper gives some numerical examples demonstrating the surprisingly strong implications for economic policy of the distinction between additive and multiplicative disutility damages.

1 Introduction

The economics and science of climate change are characterized in practically every major dimension by deep structural uncertainties. One prime source of structural uncertainty in the economic component of climate change concerns the appropriate way to represent damages from global warming. The "damages function" is a notoriously weak link in the economics of climate change, because it is difficult to specify a priori and because the results from cost-benefit analysis (CBA) or an integrated assessment model (IAM) can be very sensitive to its functional form – particularly for high temperatures. Ideally, one wants an

*Department of Economics, Harvard University (mweitzman@harvard.edu). Without tying them to the contents of this paper, I am grateful for useful critical comments by Frank Ackerman, William Nordhaus, Robert Pindyck, and Richard Tol.

analytically tractable form that captures adequately the economic reality of global warming. The existing literature offers sparse theoretical guidance and little empirical evidence about why one form or another of a damages specification should be favored. This is especially problematic for extrapolating the hypothesized welfare impacts of high temperatures far outside the ordinary range of existing data or empirical studies.

This paper attempts to introduce some structure into an exasperating situation by postulating two reasonable conditions that a reduced-form utility function of productive capacity and temperature change might plausibly obey. It turns out that these two axioms imply a utility function encompassing as special cases both the standard multiplicative form, which has traditionally been used in practice for some time now, and also an additive analogue that has appeared more recently in the literature. After some numerical exercises, the paper discusses a few of the surprisingly strong implications for economic policy of the seemingly arcane theoretical distinction between additive and multiplicative disutility damages.

2 What is the Appropriate "Damages Function"?

Right from the beginning, the specification of damages from climate change presents severe conceptual challenges. I think the basic problem is that with global warming the core welfare-related concepts are all so intertwined that it is difficult to disentangle one of them from the others. What follows in this section is an attempt from first principles to put some "practical" theoretical structure on the elusive interconnections among productive capacity, temperature change, utility, consumption, damages, and environmental amenities. The emphasis here is on the word *practical*. The purpose of this paper is to use economic theory to place some rough guidelines on functional forms that might be useful for low-probability high-impact applications. When temperature changes are modest, the functional form of a "damages function" does not much matter for the economic analysis of climate change. But with extreme temperatures well outside the range of normal experience the specification can matter a lot, and one is unfortunately left with little more than a combination of axiomatic theory, thought experiments, and numerical examples.

Throughout this paper, T stands for exogenously given global average surface temperature change above the pre-warming level. It is assumed that T can serve as an aggregate proxy for climate change.

If C is consumption *without* global warming (i.e., for $T = 0$), the relevant utility function is presumed to be of the constant relative risk aversion (CRRA) form

$$V(C) = -C^{1-\eta} \tag{1}$$

(up to a positive affine transformation), where $\eta > 1$ is the constant coefficient of CRRA. In what follows, the base case will be $\eta = 2$, which corresponds to a not unreasonable value that has already been used in the economic modeling of climate change.

Let K stand for the aggregate productive capacity of the economy. K represents *potential* consumption in the complete absence of climate change because it is defined to be what consumption *would be* without any global warming. One interpretation is that K stands for comprehensive capital, very broadly defined to include human capital, stocks of knowledge, and so forth. Slightly more specifically, when $T = 0$ the production function for consumption is $C = AK$, with units normalized so that $A = 1$ and, initially, $K_0 = 1$. For simplicity, in this paper K is taken as exogenously given at any time and exponentially growing over time at some fixed rate. The causality in the system is that productive capacity K and temperature change T are both simplistically treated as exogenously given primitives, which together determine cardinal utility U as a function of K and T, written

$$U = U(K, T), \tag{2}$$

where, for all $K > 0$ and $T > 0$, $U_1 > 0$, $U_{11} < 0$, $U_2 < 0$, $U_{22} < 0$.

It is essential to understand that $U(K, T)$ represents U only as a *reduced form* in K and T. Expression (2) does not imply that temperature (or capital) enters the utility function *directly*. The indirect pathway via which T influences U for given K will be exposited later in some detail. The purpose of this paper is to attempt an end run around mechanical curve-fitting approaches by focusing directly on some plausible risk-aversion properties of $U(K, T)$ that generalize existing formulations already being used in the literature.

What I will call for concreteness the *prototype multiplicative* utility specification combines a coefficient of relative risk aversion $\eta = 2$ in (1) with a multiplicative-quadratic damages function in the overall form (up to an affine transformation) of

$$U_m(K, T) = -\left[\left(\frac{1}{K}\right) \times \left(1 + \alpha T^2\right)\right], \tag{3}$$

where α is a positive coefficient, typically calibrated to some postulated loss of consumption, say $\approx 2\%$ for $T \approx 2°$C.

Perhaps the prime production activity motivating the "prototype multiplicative" form (3) is agriculture. In the agriculture application, temperature affects production directly by physically changing the amount produced. If there were no global warming, the production function would be $C = K$. In the presence of temperature change $T > 0$, the ability to produce consumption C out of capacity K is physically diminished by the "damages function" $1/(1+\alpha\, T^2)$, so that $C = K/(1+\alpha T^2)$ When $\eta = 2$, the relevant reduced-form utility function from (1) is $U_m(K,T) = -\left[K/\left(1+\alpha T^2\right)\right]^{-1}$, which is the same as (3). This physical image of a "damages function" of temperature multiplicatively diminishing the ability to produce consumption, writ large, motivates the popular "prototype multiplicative" approach to modeling the welfare impacts of climate change. In a context of climate change, however, this physical production image might be inaccurate because of the possibility that it can only be extended to a limited subset of other welfare-producing activities or commodities, most of which might better be described by introducing some notion of the utility of an "environmental amenity." As will be shown later, the problem of a possibly inaccurate representation of damages by the "prototype multiplicative" form might be greatly exacerbated with higher temperatures.

A *prototype additive* analogue of (3) is the utility function (up to an affine transformation)

$$U_a(K,T) = -\left[\left(\frac{1}{K}\right) + \left(1+\alpha T^2\right)\right], \tag{4}$$

where, again, α is a positive coefficient, typically calibrated to some postulated loss of consumption, say $\approx$ 2% for $T \approx 2°$C. Formulation (4) features a standard quadratic loss function, perhaps representing the diminishment of utility caused by a decline in the "environmental amenity." Note that (3) and (4) are symmetric, with the only difference being the "$\times$" sign in (3) and the "+" sign in (4). I think it is fair to say that it is hard to argue strongly for one form over the other from any basic principles.

Although they are different (if symmetric), both the multiplicative utility function (3) and the additive utility function (4) are of the constant relative risk aversion (CRRA) class. For all $K > 0$ and $T > 0$, they both share the same constant coefficient of relative K-risk aversion, denoted η (where here $\eta = 2$). And they both share the same constant coefficient of relative T-risk aversion for all K and T, denoted γ (where here $\gamma = 1$). I now derive the most general form of a utility function that simultaneously exhibits both of these characteristics for any given values of $\eta > 1$ and $\gamma > 0$ (not just for $\eta = 2$ and $\gamma = 1$). The purpose of this exercise is to embed the multiplicative and additive forms (3) and (4) within some common axiomatic framework, the better to understand how they fit, separately and together, into the somewhat confusing picture of an appropriate "damages function" for global warming.

The following two axioms generalize the basic properties of (3) and (4) into what I would consider to be a pragmatic compromise between generality and analytical tractability.

Abusing terminology somewhat, let the "coefficient of relative *consumption-risk* aversion" be given by the formula

$$\eta(K,T) = -\frac{K\, U_{11}(K,T)}{U_1(K,T)}. \tag{5}$$

Axiom 1 *For all $K > 0$ and $T \geq 0$, the coefficient of relative* ***consumption-risk*** *aversion is the constant $\eta > 1$.*

Turning now to T and the second axiom, I would say that there is just as strong an argument from first principles that the representative agent is averse to uncertainty about temperature change as that he or she is averse to uncertainty about capacity to produce consumption. For example, a representative agent might be indifferent between a sure temperature change of T=3.5°C and a 50-50 chance of T=0°C or T=5°C (whose expected value is $E[T]$=2.5°C). As the paper will show, if coefficients of relative risk aversion are constant (although different) for all K and T, then it will lead to an cardinal utility function of an analytically tractable form.

Define the "coefficient of relative *temperature-risk* aversion" as

$$\gamma(K, T) \equiv \frac{T\, U_{22}(K, T)}{U_2(K, T)}. \tag{6}$$

A second axiom, to be explained presently, is a less familiar postulate than Axiom 1. I argue (by analogy with Axiom 1) that this second axiom is a sufficiently reasonable approximation to be useful in practical applications where one wants an analytically tractable functional form that does not excessively undermine reality. The following postulate requires that the coefficient of relative *temperature-risk* aversion is the same for all levels of the economy's capacity to produce consumption and for all temperature changes. This postulate represents what I believe, in a context of global warming, is a useful combination of generality with analytical tractability.

Axiom 2 *For all $K > 0$ and $T > 0$, the coefficient of relative* ***temperature-risk*** *aversion is the constant $\gamma > 0$.*

Suppose, as in a previous numerical example, that the representative agent is indifferent between a sure temperature change of T=3.5°C and a 50-50 chance of $\widetilde{T}$=0°C or $\widetilde{T}$=5°C (whose expected value is $E[\widetilde{T}]$=2.5°C). Then an implication of Axiom 2 is that the same agent should also be indifferent between a sure temperature change of T=1.75°C and a 50-50 chance of $\widetilde{T}$=0°C or $\widetilde{T}$=2.5°C – and also between a sure temperature change of T=7°C and a 50-50 chance of T=0°C or T=10°C. A further implication of Axiom 2 is that the above type of indifference relations should hold irrespective of the level of productive capacity K. Postulating such kind of properties (which both (3) and (4) obey anyway) seems to me an acceptable price to pay in terms of loss of generality for buying into what will turn out to be a considerable degree of analytical simplicity. (It is readily confirmed that the above particular numerical example, which seem plausible to me, imply that $\gamma = .99$. The base case for the numerical examples to follow will be $\gamma = 1$, which implies that all losses will be expressed in terms of the quadratic-polynomial form T^2.)

Axiom 1 implies that

$$U(K,T) = a(T) + b(T)\, K^{1-\eta}, \tag{7}$$

where $a(T)$ and $b(T)$ are functions of T representing an affine transformation of $K^{1-\eta}$.

Axiom 2 implies that

$$U(K,T) = \alpha(K) + \beta(K)\, T^{1+\gamma}, \tag{8}$$

where $\alpha(K)$ and $\beta(K)$ are functions of K representing an affine transformation of $T^{1+\gamma}$.

It is not difficult to show that the most general utility function satisfying (7) and (8) must (up to an affine transformation) be of the functional form

$$U(K,T) = -K^{1-\eta} - \alpha_A\, T^{1+\gamma} - \alpha_M\, K^{1-\eta}\, T^{1+\gamma}, \tag{9}$$

where α_A and α_M are non-negative constants, at least one of which must be positive.

It is readily confirmed that the specification (9) has what might be considered intuitively-required essential properties for a cardinal utility function $U(K,T)$. When $T > 0$ and $K > 0$, note that

$$\frac{\partial U}{\partial K} > 0,\ \frac{\partial^2 U}{\partial K^2} < 0,\ \frac{\partial U}{\partial T} < 0,\ \frac{\partial^2 U}{\partial T^2} < 0, \tag{10}$$

while, when $T = 0$ and $K > 0$, note that $U_2(K,0) = 0$ and also that $U(K,0) = V(K) = -K^{1-\eta}$, which condition confirms that productive capacity represents what consumption would be without global warming.

Define the "damages function" $D(K,T)$ to be the implicit solution of

$$V((1 - D(K,T))\, K) = U(K,T), \tag{11}$$

with $V(\bullet)$ defined by (1)

Then plugging (1) and (9) into (11) yields the formula

$$D(K,T) = 1 - \left[K^{1-\eta} + \alpha_A\, K^{1-\eta}\, T^{1+\gamma} + \alpha_M\, T^{1+\gamma}\right]^{-\frac{1}{\eta-1}}. \tag{12}$$

The "prototype multiplicative" form (3) is a special case of (9) for parameter values $\eta = 2$, $\gamma = 1$, $\alpha_A = 0$, $\alpha_M > 0$, and therefore it obeys Axioms 1 and 2. Plugging these parameter values into (12), the utility function (3) is equivalent to what I will call in this paper the "prototype multiplicative" damages function

$$D_M(K,T) = \frac{\alpha_M \, T^2}{1 + \alpha_M \, T^2}, \tag{13}$$

which actually holds more generally for *any* value of $\eta > 1$ (when combined with parameter values $\alpha_A = 0$, $\alpha_M > 0$, and $\gamma = 1$).

The "prototype additive" form (4) is a special case of (9) for $\alpha_M = 0$, $\alpha_A > 0$, $\gamma = 1$, $\eta = 2$, and therefore obeys Axioms 1 and 2. Plugging these parameter values into (12), the utility function (4) is equivalent to what I will call in this paper the "prototype additive" damages function

$$D_A(K,T) = \frac{\alpha_A \, K \, T^2}{1 + \alpha_A \, K \, T^2}, \tag{14}$$

where (in a slight abuse of terminology) I call (14) "prototype additive" because of its connection to the additively separable utility function (4).

Both the "prototype multiplicative" formulation (13) and the "prototype additive" formulation (14) accept that utility can be reduced to the one-variable functional form $U = V((1-D)\,K)$, and both specifications model the impact of higher temperature changes as entering the utility function via higher values of damages D. The issue is *not* that reduced-form utility can be written as $U(K,T) = V((1 - D(K,T))\,K)$ for the "prototype multiplicative" specification (13), but that reduced-form utility cannot be so expressed for the "prototype additive" specification (14). In both specifications, damages already include the impact of temperature change and enter the utility function via the nested form $V((1 - D(K,T))\,K)$. The real issue is the different *values* of $D(K,T)$ given by (13) and (14).

Conditional upon the standard fallback specification $\eta = 2$ and $\gamma = 1$, for the more general situation than (13) and (14) where α_M and α_A might both be positive the damages function from (12) becomes

$$D(K,T) = \frac{\alpha_A \, K \, T^2 + \alpha_M \, T^2}{1 + \alpha_A \, K \, T^2 + \alpha_M \, T^2}, \tag{15}$$

which nests (13) and (14) as special cases.

The standard way of conceptualizing a damages function in the literature is that it represents the effects of temperature via something like an actual multiplicative physical change in the material output being produced, as with the simple agricultural examples sometimes used to illustrate (3). However, tying welfare damages to this physical image of a temperature change that multiplicatively shifts the production function is not essential. To see this, define abstractly an aggregate "environmental amenity" good E by the expression

$$E(T,K) \equiv \frac{1}{1+\alpha_A T^2 + \alpha_M T^2/K}, \tag{16}$$

along with the symmetrically-additive utility function

$$u(K,E) \equiv -\left[\frac{1}{K}+\frac{1}{E}\right]. \tag{17}$$

Simultaneously combine (16) with (17), and combine (15) with (11). Then, after a few algebraic manipulations, note that

$$u(K,E(T,K)) = V((1-D(K,T))\,K)+1, \tag{18}$$

with $D(K,T)$ defined by (15).

The significance of (18) is the following. One way to interpret or to conceptualize the impact of temperature changes on welfare is via a channel of multiplicative damages to productive capacity in the form of (11) coupled with the damages function (15). An alternative way to interpret or to conceptualize this impact is via a channel whereby, when there are higher temperatures, then there is less of the "environmental amenity" good (16), which lowers utility in (17). Equation (18) signifies that, at least theoretically, the damages-function approach and the environmental-amenity approach are isomorphic within the above setup. For every damages function (15), there is an environmental-amenity good (16) that gives the same results – and for every environmental-amenity good (16), there is a damages function (15) that gives the same results. In principle, it is irrelevant to an economic analysis of global warming here whether the welfare impacts of increased temperatures are envisioned as going through the damages-function channel, or through the environmental-amenity channel, or through some combination of the two.

So far the discussion has been mainly theoretical. The next section explores some numerical examples. The idea is to get a ballpark quantitative sense of the differing policy implications of a distinction between "prototype multiplicative" and "prototype additive" global-warming damages.

3 What Difference Might This Make Empirically?

The main application of the theoretical result of this paper shows that CRRA in K and CRRA in T, with CRRA coefficients $\eta = 2$ and $\gamma = 1$, simultaneously *implies* and *is implied* *by* the relatively simple damages function (15). This equivalence result is a useful tool for narrowing the realm of discourse down to a discussion of the values of the two remaining parameters α_M and α_A within a very specific functional form. But this is about as far as one can go with pure theory alone. The rest depends on some judgements about the relative values of the parameters α_M and α_A. It is essentially impossible to discriminate empirically between the forms (13) and (14) on the basis of any known data or existing studies, so that readers must subjectively decide for themselves which parameter values they feel most comfortable with when extrapolating (9) and (15) for higher values of T and K. In the numerical exercises that follow, I concentrate on comparing the economic implications of the pure "prototype multiplicative" damages specification (13) with the economic implications of the pure "prototype additive" damages specification (14) (both of which correspond to CRRA coefficients $\eta = 2$ and $\gamma = 1$).

When current $K = K_0$ is normalized to unity, then $D_M = D_A$ for all T, and the same calibration can be used to fix the same value of $\alpha = \alpha_M = \alpha_A$ in both cases. For the numerical exercises that follow, I calibrate α so that 2% of current welfare-equivalent consumption would be lost if current temperature change were T=2°C.[1] Then α satisfies the equation

$$U(.98, 0) = U(1, 2), \tag{19}$$

where U is either U_M or U_A, both of which obey (19) and deliver the same value of α. The solution of equation (19) is $\alpha = .0051$, which is the base-case value used in all of the following numerical examples.

When $K \approx K_0 = 1$, there is not much difference between (13) and (14), but for large values of $K \gg 1$ and $T \gg 0$ the distinction becomes significant because, from comparing (13) with (14), D_A is then substantially higher than D_M. For any *given* $T > 0$ in formula (13), damages D_M are always the same constant, irrespective of the level of K. This property, that damages D_M are independent of K for given T, is a hallmark of the "prototype multiplicative" specification (13). It might seem like a somewhat peculiar property because one might think that in a rich high-K high-C world the fraction of productive capacity that people would willingly sacrifice to avoid altogether any temperature change, D, might be higher than in a poor low-K low-C world. Note that D_A in (14) has just this desirable characteristic.

Let time be denoted t, with the present time normalized to $t = 0$. Thus, capacity to produce consumption at future time t is $K(t)$, while temperature change at time t is $T(t)$. For convenience, the chosen normalization is $K(0) = 1$ and $T(0) = 0$. To compare and con-

[1] Such type of calibration is performed in Nordhaus (2008), Sterner and Persson (2008), and many other studies.

trast in familiar language the basic empirical consequences of the "prototype multiplicative" damages function (13) with the "prototype additive" damages function (14), I now ask the following question. What is the total willingness to pay as a fraction of *current* consumption (here $C(0) = K(0) = 1$) that the representative agent would accept to eliminate the temperature $T(t) > 0$ at time $t > 0$ by reducing it all the way down to $T(t) = 0$? Call this value ω. Suppose that the rate of pure time preference or "utility discount rate" is δ. Then ω must satisfy the equation

$$U(1,0) - U((1-\omega),0) = \exp(-\delta t)\left[U(K(t),0) - U(K(t),T(t))\right], \tag{20}$$

where U in this numerical application is either U_M or U_A.

Suppose that capital stock, representing the potential capacity to produce consumption, grows at rate g, so that $K(t) = \exp(gt)$. Then plugging (3) and (4) respectively into (20), after some algebraic rearranging one obtains

$$\omega_M = \frac{\alpha T(t)^2}{\exp((g+\delta)t) + \alpha T(t)^2} \tag{21}$$

and

$$\omega_A = \frac{\alpha\, T(t)^2}{\exp(\delta t) +\, \alpha\, T(t)^2}. \tag{22}$$

The difference between the "prototype multiplicative" willingness to pay (21) and the "prototype additive" willingness to pay (22) comes down to the latter being free of the powerful dampening term $\exp(gt)$ in the denominator. To give a numerical example emphasizing the significance of this kind of distinction, suppose that $g = 2\%$ and $\delta = .5\%$ (along with $\eta = 2$ and $\gamma = 1$). The famous Ramsey interest-rate formula is $r = \delta + \eta g$, which implies that the above parameter values determine a real interest rate of $r = 4.5\%$. This is arguably a high discount rate to use in CBAs of climate change – and in that sense $r = 4.5\%$ biases my numerical results here against the case I am trying to make that high future temperatures can "bite" now. For the above parameter values, I calculate the current total willingness to pay to avoid altogether a hypothetical "business as usual" temperature change of 5°C a century from now.[2]

Plugging $g = 2\%$ and $\delta = .5\%$ into (21) and (22), I now ask the following question. What is the total willingness to pay (at time $t = 0$, as a fraction of $C(0) = K(0) = 1$) to avoid altogether $T(100) = 5°C$ by reducing it to $T(100) = 0°C$? The answer under the

[2] As partial justification for $T(100) = 5°C$, a recent comprehensive MIT study (Socolov et al, 2009) estimated a median "business as usual" temperature change for 2100 (91 years from now) at 5.2°C. Many other studies could be cited to confirm that $T(100) = 5°C$ is within about one or so standard deviations of what is possible.

multiplicative specification is $\omega_M = 1\%$, while under the additive specification it is $\omega_A = 7.2\%$. This difference is substantial and might be interpreted as giving sharply contrasting policy advice about the urgency of climate-change mitigation efforts. Another way to envision this dramatic difference comes from asking the following question. How much of a temperature change a century from now would justify spending 7.2% of current consumption to eliminate it altogether? With additive utility (14), the answer from (22) was shown to be $T(100) = 5°\text{C}$. With multiplicative utility (13), the answer from inverting (21) is $T(100) = 14°\text{C}$! Willingness to pay to avoid completely all global warming represents an unrealistic extreme measure, but it does give some very rough sense of the relative magnitudes of the potential benefits of climate-change mitigation under different damages functions.

The reader can plug in other parameter values and perform other numerical experiments, but I think the empirical conclusions are likely to be in the spirit of what emerges from the above numerical exercises. In this sense it might be argued that, relative to the multiplicative formulation (13), the additive formulation (14) does not nearly so significantly blunt the current welfare impacts of large future temperature changes. When K and T are large, the "prototype additive" specification (14) makes it much harder to substitute lower consumption for higher temperatures than the "prototype multiplicative" specification (13). Such a conclusion represents a more or less intuitive consequence of the idea that marginal rates of substitution should generally be more sensitive to changes in temperature for additive utility than for multiplicative utility.

One lesson to be drawn from these simple numerical examples is that a seemingly arcane distinction between an additive and a multiplicative interaction of temperature change with productive capacity might have big empirical consequences. I think that this is the overarching message of the paper, and that it transcends theoretical debates about multiplicative vs. additive utility functions. If tradeoffs depend critically upon the functional forms by which temperature changes enter production or utility functions, then it becomes yet another example of structural uncertainty exerting a decisive influence on climate-change policy. (Here the "structural uncertainty" concerns the specification of temperature damages.)

In related work, Sterner and Persson (2008) introduced into the damages-function debate an important numerical exercise with a constant elasticity of substitution (CES) utility function that can be interpreted within the framework of this paper as having the form

$$W(K,T) = \frac{1}{1-\eta}\left[(1-b)K^{\frac{\sigma-1}{\sigma}} + bE^{\frac{\sigma-1}{\sigma}}\right]^{\frac{(1-\eta)\sigma}{\sigma-1}}, \tag{23}$$

where the environmental good E is defined as $E \equiv 1/(1+\alpha T^2)$, and σ is the elasticity of substitution between K and E.

The base-case parameter values chosen by Sterner and Persson are $\sigma = \frac{1}{2}$ and $\eta = 2$. For these parameter values, it is readily seen that (23) with $E \equiv 1/(1 + \alpha T^2)$ equals (17), (16) with $\alpha_M = 0$, and that both expressions are equivalent to the additive specification (4), (14) (up to an affine transformation). Empirically then, in this paper my "prototype additive" specification (4), (14) can piggyback on the the numerical findings of the Sterner and Persson study. They found that plugging their CES utility function (23) with $E \equiv 1/(1 + \alpha T^2)$ (equivalently, my "prototype additive" specification (4), (14)) into William Nordhaus's pioneering DICE model[3] yields a significantly more stringent emissions policy than Nordhaus found for his "prototype multiplicative" form (3), (13). With Nordhaus's "prototype multiplicative" specification (3), (13), optimal CO_2 concentrations in DICE reach almost 700 ppm by 2150 and the optimal average world temperature eventually increases by more than 4°C. In Sterner and Persson's CES specification (23) with $E \equiv 1/(1 + \alpha T^2)$ (equivalent to my "prototype additive" specification (4), (17), optimal CO_2 concentrations are stabilized at under 450 ppm by the end of this century and the optimal global mean temperature change stays below 2°C. I think the critical issue here is that the "prototype multiplicative" specification (3), (13) used in DICE and many other IAMs makes it very easy to substitute consumption for temperatures because the implicit elasticity of substitution between K and $E \equiv 1/(1 + \alpha T^2)$ is $\sigma = 1$.

As an empirical matter then, the study of Sterner and Persson is consistent with the numerical findings of this paper that a seemingly obscure distinction between multiplicative and additive interactions of productive capacity with temperature change makes a significant difference for climate-change policy. The underlying reason should be more or less apparent from previous discussions in this paper. Furthermore, the significant empirical difference found by Sterner and Persson emerges from a deterministic version of the DICE model (no numerical simulations of probability distributions), along with a relatively high rate of pure time preference $\delta \approx 1.5\%$ per year. If one introduces fat-tailed climate change uncertainty, along with even tiny subjective probabilities of low rates of pure time preference δ, the difference between additive and multiplicative combining of K and T becomes an overwhelmingly dominant force in determining optimal climate-change policies.[4] Modelers using the "prototype multiplicative" formulation (3), (13) are sometimes unaware of the possible sensitivity of their model's outcomes to this particular specification.

[3] See Nordhaus (2008).

[4] This claim is detailed in Weitzman (2009).

4 Concluding Comments

Issues of deep structural uncertainty are fundamental to any economic analysis of climate change. This paper shows how structural uncertainty concerning the appropriate form of high-temperature damages might greatly influence, to the point of almost predetermining, the outcome of climate-change CBAs and IAMs. The moderate "policy ramp" that emerges from standard CBAs and IAMs of climate change may, at least partially, be an artifact of the high-substitution "prototype multiplicative" utility function that is routinely used in standard CBAs and IAMs.

This paper relies on two plausible axioms that postulate CRRA in capacity to produce consumption and CRRA in temperatures. The basic result of the paper shows that these two axioms reduce a general discussion about the appropriate "damages function" into a specific discussion about choosing the parameters of a parsimonious analytically-tractable functional form that includes the "prototype multiplicative" and "prototype additive" specifications as special cases. Empirically, the numerical examples cited in this paper indicate that the "prototype additive" form favors a far more stringent emissions policy than the gradualist policy ramp of the "prototype multiplicative" form.

Plausible axioms can be extremely useful for narrowing down a universe of possible specifications into a particular functional form that can be of practical use for weighing tradeoffs and making decisions – but the axioms rarely give a decisive final word. No matter how theoretical debates about multiplicative vs. additive damages are eventually resolved, fragility of policy to postulated forms of utility functions or damages functions is an unsettling empirical finding for the economics of climate change. With this kind of fundamental non-robustness, the outcomes of CBAs or IAMs are held hostage to core structural uncertainties concerning how high temperature change and high productive capacity should be combined to yield utility. Such a dismal message is not intended to cause despair for the economics of climate change, nor to negate the need for further study and numerical simulations to guide policy. Instead, this message is just another warning, in a growing series of cautionary tales, that the *particular* application of CBAs or IAMs to climate change seems more inherently prone to being dependent on subjective judgements about structural uncertainties than most other, more ordinary, applications of CBAs or IAMs.

References

[1] Nordhaus, William. *A Question of Balance.* Yale University Press, 2008.

[2] A.P. Sokolov, P.H. Stone, C.E. Forest, R. Prinn, M.C. Sarofim, M. Webster, S. Paltsev, C.A. Schlosser, D. Kicklighter, S. Dutkiewicz, J. Reilly, C. Wang, B Felzer, H.D. Jacoby. "Probabilistic forecast for 21st century climate based on uncertainties in emissions (without policy) and climate parameters." *Journal of Climate*, preprint, 2009: http://ams.allenpress.com/perlserv/?request=get-abstract&doi=10.1175%2F2009JCLI2863.1

[3] Sterner, Thomas, and U. Martin Persson. "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate." *Review of Environmental Economics and Policy*, 2008 (Winter), 2(1), pp. 61-76.

[4] Weitzman, Martin L. (2009). "Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting." Mimeo, May 28, 2009, posted on the author's web page and also available from the author upon request.