Should the climate tail wag the policy dog?

Gerard H. Roe

Department of Earth and Space Sciences,

University of Washington,Seattle, WA.

and

Yoram Bauman

Program on the Environment,

University of Washington, Seattle, WA.

October 25, 2010

Abstract

The small but stubbornly unyielding possibility of a very large long-term response of global temperature to increases in atmospheric carbon dioxide can be termed the fat tail of high climate sensitivity. It has been suggested that the fat tail should properly dominate a rational policy strategy, if the damages associated with such high temperatures are large enough. Here we explore the role of two robust physical properties of the climate system in this argument: the enormous thermal inertia of the ocean, and the long timescales associated with high climate sensitivity. If climate sensitivity in fact proves to be high, these two properties prevent the high temperatures in the fat tail from being reached for many centuries. An economic analysis that assigns large damages to high temperatures shows that the impact on welfare-equivalent consumption would only exceed a percentage point for the case of very strong climate forcing (a quadrupling of carbon dioxide) combined with very low long-term growth rates (less than one percent). The same strong limits also govern the conditions under which a reactive damage function, that assumes magnified damages at higher temperatures, matters. We suggest that both the assumptions within the economic framework, and the reliance on global mean temperature as a metric for assessing climate damage, can be questioned.

1 Introduction - the fat tail of climate sensitivity

Climate sensitivity – the long-term response of global-mean, annual-mean surface temperature to a doubling of carbon dioxide above pre-industrial values – has long been a benchmark by which to compare different estimates of the planet's climatic response to changes in radiative forcing. A doubling of carbon dioxide increases the radiative forcing by about $\Delta R_{2\times} \approx 4$ Wm^{-2}. In a remarkable piece of analysis, Svante Arrhenius (1896) made the first quantitative estimate that, in response, the equilibrium global-mean temperature would increase by $\Delta T_{2\times} \approx 5$ °C. A major reassessment came in 1979 with the National Research Council Charney Report. Reviewing the intervening advances in science, Charney (1979) estimated $\Delta T_{2\times}$ between 1.5 and 4.5 °C (described as the 'probable error'), a range that has changed only incrementally ever since.

Modern estimates provide an answer in terms of more precise probability distributions, or probability density functions (PDFs). The Intergovernmental Panel on Climate Change (IPCC, 2007) report gives a 'likely' range (2-in-3 chance) of $\Delta T_{2\times}$ lying between 2 and 4.5 °C, and concludes it is 'very unlikely' (< 1-in-10) to be less than 1.5 °C. Note that this summary is consistent with a peculiarity of a large number of other studies, namely their inability to rule out the small possibility of $\Delta T_{2\times}$ being very much larger than the canonical 1.5 to 4.5 °C range. One of the important achievements in recent climate science has been to establish great confidence in the lower bound on climate sensitivity, but the upper bound has proven much less tractable (e.g., Knutti and Hegerl, 2008). Allen et al. (2006) have argued that this fat tail of possibly-high climate sensitivity is a fundamental feature of $\Delta T_{2\times}$ estimates from observations, and Roe and Baker (2007) have argued that it is a fundamental feature of $\Delta T_{2\times}$ estimates from numerical climate models. It will require improbably large reductions in uncertainties about the radiative forcing the planet has experienced—or, equivalently, in our uncertainty about physical feedbacks in the climate system—to substantially remove the skewness. Some studies (e.g., Annan and Hargreaves, 2006) have combined multiple estimates in a Bayesian framework and argued this can yield narrower and less skewed distributions for λ. However the answers derived from such an approach are critically sensitive to how independent the different estimates are, and what the Bayesian prior assumptions are. Both factors are fiendishly elusive to pin down objectively. Knutti and Hegerl (2008), and Knutti et al., (2010) provides a good review of the scientific issues involved. Efforts continue in this direction (Annan and Hargreaves, 2010), but overall, it seems prudent to assume that estimates of $\Delta T_{2\times}$ will not change substantially for the foreseeable future.

Several macroeconomic analyses of the costs of climate change, and of the correct willingness-to-pay to avoid it, have argued that there is a sting in the fat tail of $\Delta T_{2\times}$ (e.g., Weitzman, 2009a,b, 2010). This is seen as follows: if, as is reasonable to assume, climate damages increase nonlinearly with temperature, then the tail of the PDF is weighted more strongly that the middle of the PDF in calculations of the expected damages. Moreover, if damages are highly nonlinear with temperature, then even the very smallest chance of absolutely apocalyptic consequences might properly dominate a rational policy. A close analogy is the St Petersburg paradox, a coin-flipping wager with an expected value dominated by low-probability, high-valued outcomes.

The extent to which this is a worry depends in part on how rapidly the global mean temperature rises towards its equilibrium value $\Delta T_{2\times}$. The focus of the present paper is to emphasize that there are some very strong physical constraints on how quickly the climate system can warm up, and that these same physical constraints also control how possible future climate trajectories evolve in time. The timescales involved play a critical role in economic analyses that include discounting of future damages.

2 The transient evolution of possible future climates

A key player in the physical system is the enormous thermal inertia represented by the deep ocean. The whole climate system cannot reach a new equilibrium until the deep ocean has also reached equilibrium. In response to a positive climate forcing (i.e., a warming tendency), the deep ocean draws heat away from the surface ocean, and so buffers the surface temperature changes, making them less than they would otherwise be. The deep ocean is capable of absorbing enormous amounts of heat, and not until this reservoir has been exhausted can the surface temperatures attain their full, equilibrium values.

A second key player is the inherent relationship between feedbacks and adjustment time scales in physical systems. If it transpires that we do, in fact, live on a planet with a high climate sensitivity, it will be because we live on a planet with strong positive feedbacks. In other words, the net effect of all of the dynamic processes (clouds, water vapor, ice reflectivity, etc.) is to strongly amplify the planet's response to radiative forcing. In this event, it would mean that we live on a planet that is inefficient in eliminating energy perturbations: a positive feedback reflects a tendency to retain energy within the system, inhibiting its ultimate emission to space, and therefore requiring a larger temperature response in order to achieve energy equilibrium. Moreover, it is generally true that, all else being equal, an inefficient system takes longer to adjust than an efficient one. A useful rule-of-thumb is that the relevant response time of the climate system is given by the effective thermal inertia of the deep ocean multiplied by the climate sensitivity parameter (defined as $\Delta T_{2\times}/\Delta R_{2\times}$, see, e.g., Roe, 2009). This behavior is absolutely fundamental and widely appreciated (e.g., Hansen et al., 1985; Wigley and Schlessinger, 1985). In the context of the PDF of climate sensitivity, its effects have been reviewed in Baker and Roe (2009).[1]

[1] In other contexts, the timescale of the ocean response in sometimes cited as being about 1500 yrs, based on the

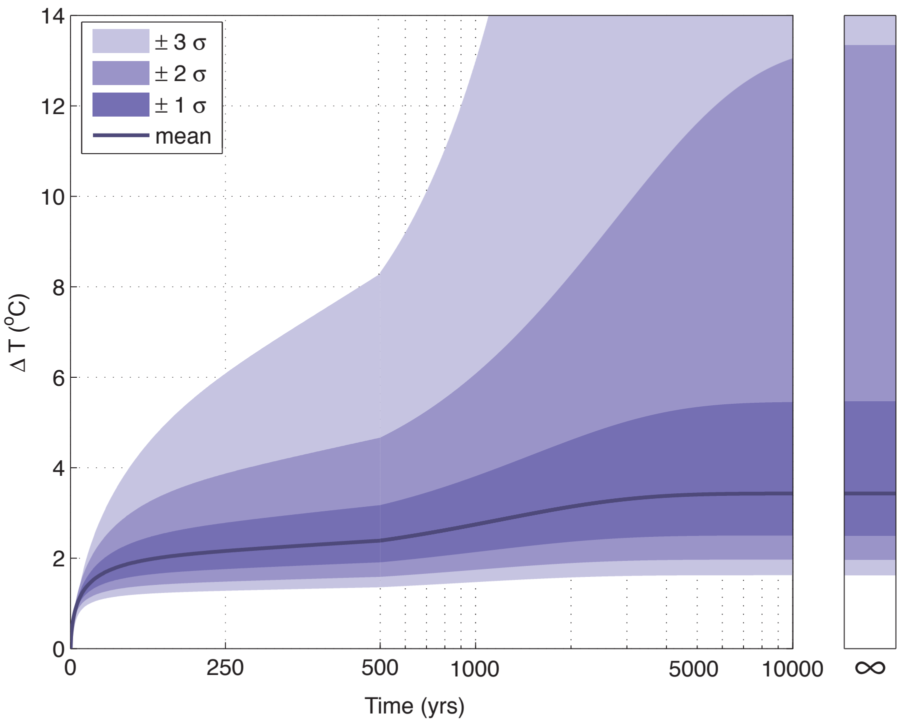

Figure 1: It takes a very long time to realize the full equilibrium PDF of climate sensitivity. Instantaneous doubling of CO2, standard params. Note 1,2,3 σ ranges encompass 68.3, 95.5, and 99.7% of possibilities, respectively. See text for further explanation of the uncertainties and see Figure 2 for the full shape of the distribution.

The dramatically long timescale for the development of the fat tail is shown in Figure 1. It presents the evolution in time of the envelope of possible climate trajectories of global-mean temperature, for an instantaneous doubling of carbon dioxide. Note the change to a logarithmic time axis after 500 years. The shading in Figure 1 represents the one-, two-, and three-standard deviation ranges for climate sensitivity, encompassing 68.3, 95.5, and 99.7% of possibilities, respectively. The figure is generated using a simple climate model which represents both atmospheric feedbacks and the uptake of heat by both the upper and deep ocean. The equations and parameters follow those of Baker and Roe (2009), who match best current estimates of heat uptake and climate feedbacks. The model is virtually identical to a host of other equivalent models, which have been shown to be fully capable of emulating more complete numerical models, and also historical observations at the global scale (e.g., IPCC, 2007). Such models are regularly used to make long-term climate predictions (e.g., IPCC, 2007). Their flexibility means that parameters can be easily varied, and that uncertainties can be fully explored. Results are insensitive to plausible variations in these climate parameters.

The uncertainty ranges used are fully consistent with observed estimates of climate sensitivity (e.g., Allen et al., 2006), with the multi-thousand model experiments of the innovative *climateprediction.net* program (Stainforth et al., 2005), and with the IPCC ranges of uncertainty in $\Delta T_{2\times}$ (e.g., $\sim$ 75% of possible $\Delta T_{2\times}$s lie at less than $\sim$ 4.5 °C). At the extreme limit of the range we consider, values of climate sensitivity at 3σ would actually imply that the climate system is slightly unstable, in other words a runaway greenhouse effect.[2] Figure 1 therefore illustrates that even a

typical residence time of a water molecule. It is worth noting that this is not equivalent to the timescale discussed here, which applies to the equilibration of a heating anomaly.

[2]See Baker and Roe, 2009 for details, but we use a model estimate for the climate feedback parameter of $f = 0.65$, and a standard devtation of $\sigma = 0.13$, meaning the 3σ value is $f = 1.04$. Values of $f > 1$ means the system is formally unstable.

planet formally headed to oblivion can take a very long time to get there because of the ocean's capacity to absorb heat. The consensus of the climate science community would be that the likelihood we are actually on a runaway greenhouse trajectory is vanishingly small (e.g., Solomon et al., 2010), although the objective basis for this belief is subtler and less well-established than is commonly appreciated. It is nonetheless true that almost all climate scientists would view the uncertainties in Figure 1 to represent an overestimate. In other words, we are effectively considering a maximum possible upper bound on the probabilities residing within the fat tail.

Let $h(\Delta T)$ stand for the PDF of possible future global mean temperatures. The abiding impression from Figure 1 is that trajectories at the low end of $h(\Delta T)$ rapidly adjust to their equilibrium values over a few decades or a century, whereas those at the high end take thousand of years even to approach their equilibrium values. For the reasons given above, climate trajectories stay tightly bunched over the course of the first few centuries, diverging only slowly thereafter, and all of that divergence occurs toward higher temperatures. Baker and Roe (2009) show, at any given time, that the flux (i.e., growth) of probabilities past a given temperature, ΔT_1, is given by the product of $h(\Delta T_1)$ and the rate of warming on the trajectory passing through ΔT_1. In this light, the fat tail is walloped twice: firstly, $h(\Delta T_1)$ is small if it is in the tail; secondly, the rate of warming is generally declining with time (e.g., Figure 1). Therefore the fat tail can fill up only very gradually.

Time slices of the envelope of $h(\Delta T)$ are shown in Figure 2. This provides a feel for the evolution of the fat tail over time. At $t = 0$ the shape of $h(\Delta T)$ is actually Gaussian; it acquires skewness only gradually over time (e.g., Baker and Roe, 2009).

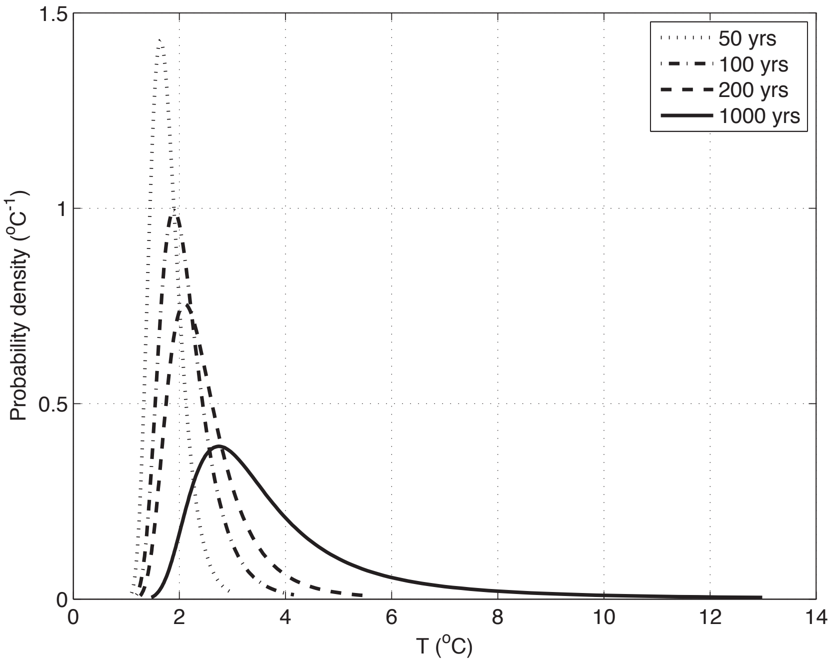

Figure 2: $h(T)$ at different times, in response to an instantaneous doubling of CO_2 at $t = 0$.

3 Costing the Earth

How does the evolution of the climate tail affect calculations of the cost of damages associated with climate change? The analysis that follows draws heavily on Weitzman (2010), which provides a simple and incisive framework for calculating the costs of avoiding or insuring against climate change. In order to explore the significance of the fat tail, Weitzman suggests comparing two different "climate damages" functions. The first is a quadratic function of global mean temperature:

$$C_Q = \frac{1}{1 + \alpha \Delta T^2}, \tag{1}$$

C_Q represents the "welfare equivalent" consumption as a fraction of what consumption would have been were $\Delta T = 0$. Weitzman takes the rather precise value of $\alpha = 2.388 \times 10^{-3}$ in order to match the damage function of Nordhaus (2008), for which $C_Q(2^{\text{o}}\text{C}) \approx 99\%$. He also considers a "reactive" damages function with an extra term:

$$C_R = \frac{1}{1 + \alpha \Delta T^2 + \beta_2 \Delta T^\gamma}. \tag{2}$$

By choosing $\beta = 5.075 \times 10^{-6}$ and $\gamma = 6.754$, Weitzman thus creates a damages function that has apocalyptic consequences for higher global-mean temperatures ($C_R(6^{\text{o}}\text{C}) = 50\%$, and $C_R(12^{\text{o}}\text{C}) = 1\%$). The functional choices and parameters are chosen purely for illustrative purposes.[3]

What should be our "willingness to pay" to avoid the climate damages associated with either C_Q or C_R, or to insure against them? Weitzman defines utility, or welfare, as a simple function of consumption, $U(C)$, and also assumes a long-term exponential economic growth rate, g. In a world experiencing climate change, consumption accelerates because of the long-term growth rate but is also retarded by the climate damages that accompany warming. Thus for every possible trajectory of future temperature, $\Delta T(t)$, there are accompanying trajectories for consumption, $C(t)$ and welfare, $U(t)$.

Weitzman then compares this world to one in which consumption is reduced by a constant fraction, $\hat{C}$, with the remainder potentially used to avoid or insure against climate changes. In this world, this fractionally reduced consumption still grows at rate g, but does not suffer climate damages. $\hat{C}$ can be found by calculating the integrated welfare of these two worlds up to a decision horizon ($\equiv t_h$, the duration of time into future that any decision is to be based on), and equating them:

[3] Note that in all cases ΔT is measured in $^{\text{o}}$C, and so the units of α and β depend on the choice of exponents. This could be avoided by dividing T by a normalizing temperature.

$$\int_0^{t_h} U(\hat{C}e^{gt})dt = \int_0^{t_h} U(C_{Q,R}(\Delta T(t))e^{gt})dt. \quad (3)$$

In principle then, $1 - \hat{C}$ ought to be the upper bound on our willingness to pay, in terms of a fraction of consumption, in order to avoid or insure against climate damages. Following reasonably standard practices, Weitzman sets $U = C^{1-\eta}/(1-\eta)$, and chooses $\eta = 3$ and $g = 2\%\ \text{yr}^{-1}$. The rate of pure time preference, ρ, is set to zero, meaning that the welfare of future generations are given equal weight to the current. In an optimizing framework, the effect of these choices is that future reductions in consumption are discounted at an exponential rate of $\eta g = 6\%\ \text{yr}^{-1}$ because future generations are richer than the current generation and therefore suffer a smaller loss in utility from reduced consumption at the margin.

Note this this rate of $6\%\ \text{yr}^{-1}$ is not quite the same as the exponential behavior in Equation (3). That exponent is $g \times (1 - \eta)$ or in other words $4\%\ \text{yr}^{-1}$. It is this rate of $4\%\ \text{yr}^{-1}$ that dominates the time-dependent behavior of the integrands, and thus is the key driver of the results. Different combinations of η, g, and ρ that still produce $4\%\ \text{yr}^{-1}$ will yield essentially the same results as those presented here. We explore later the consequences of varying the assumed growth rate.

Weitzman (2010) also takes $t_h = \infty$, which is an appropriately conservative choice, and allows for analytical solutions to his imposed climate trajectory. He also assumes a particular scenario for the consumption on the right hand side of Equation (3): $C = 1$ for a 150-year period into the future, and thereafter it is equal to either C_Q or C_R assuming that the full equilibrium PDF for global mean temperature applies. The expected value for $\hat{C}$ can be calculated. In these scenarios there is a critical difference between whether C_Q or C_R applies.

The previous section reviewed the strong physical constraints on how $h(\Delta T(t))$ evolves with time.

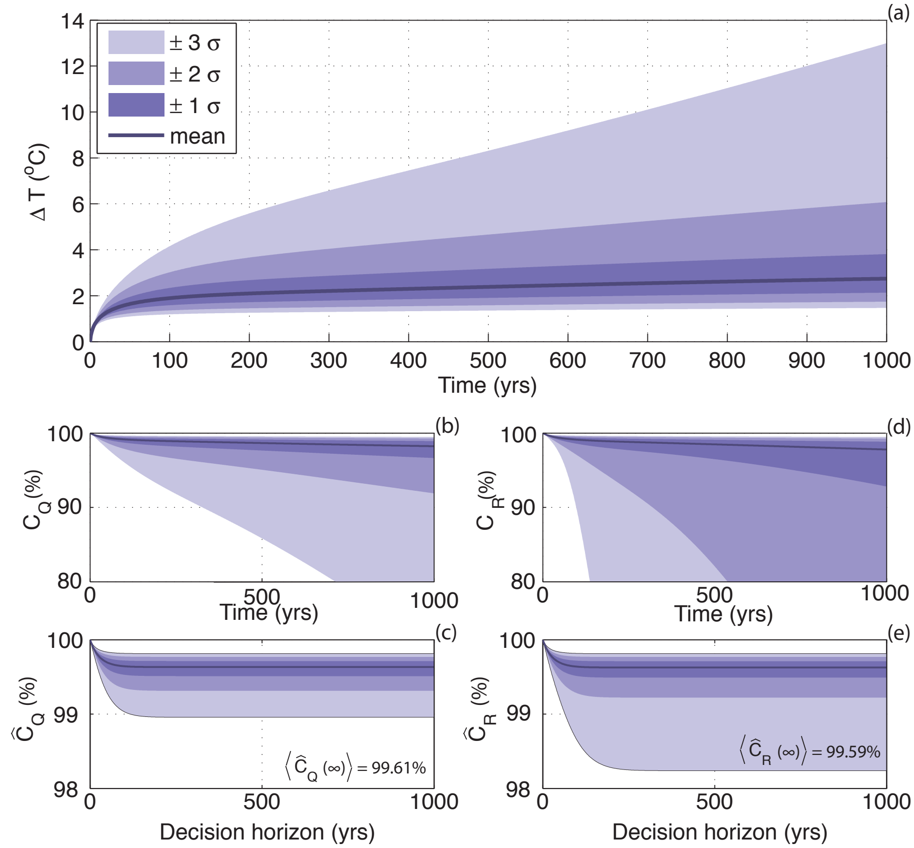

Figure 3: A closer look at the next thousand years, for an instantaneous doubling of CO_2 at $t = 0$: (a) Evolution of the PDF of possible future climates, $h(\Delta T)$; (b) evolution of the PDF of possible climate damages, using the quadratic damages function, $h(C_Q)$; (c) welfare-equivalent consumption, $h(\hat{C}_Q)$, as a function of the decision horizon; (d) and (e) are the same as panels (b) and (c), but using the reactive damages function, C_R. For all calculations, $\eta = 3$, $g = 2\%\text{yr}^{-1}$. Even for 3σ possibility of climate sensitivity, and for the reactive damages function, the willingness-to-pay remains at less than 2% equivalent consumption. The expectation values of the $\hat{C}_{Q,R}$ distributions at $t_h = \infty$ (i.e., the mean of $\hat{C}_{Q,R}$, given the PDFs) are given in angle brackets notation in the lower panels.

From Equations (1) to (3) the evolving PDFs of $C_{Q,R}$ and $\hat{C}_{Q,R}$ can also be calculated (Figure 3). Because the mode of $h(\Delta T)$ remains below 3°C over the whole millennium, there is little difference between the modes of C_Q and C_R. At the far end of the fat tail of possibilities, the 3σ trajectory reaches 6°C at around 250 years which, were it to transpire, would cause both C_Q and C_R to be a significant fraction of consumption (for instance, $C_R(3\sigma) = 47\%$ at year 250).

When cast in terms of welfare-equivalent consumption (i.e., $\hat{C}_{Q,R}$), however, we reach two important conclusions. One is that there is only a small difference between $\hat{C}_Q$ and $\hat{C}_R$, and then only at the most extreme of the future climate possibilities (above 2 σ, or about the top 5%). Both $\hat{C}_Q$ and $\hat{C}_R$ asymptote to a constant value, and change little for decision horizons beyond 150 years. As a final measure of the difference between the functions, the expectation values of the welfare-equivalent consumption at $t_h = \infty$ (i.e., the statistical means of $\hat{C}_{Q,R}(\infty)$) differ by 0.02%.

The second and arguably more important conclusion is that willingness-to-pay is under 2% for even the most extreme climate possibilities. This is largely because long periods of time are required to reach high temperatures regardless of the long-run climate sensitivity. Saying that $C_R(3\sigma) = 47\%$ at year 250 sounds impressive, but with consumption growing at 2% per year this simply means that climate change will reduce consumption from 141 times current consumption to only 53 times current consumption. Under the assumption of diminishing marginal utility of consumption it is not surprising that intertemporal optimization comes to the conclusion that current generations should be unwilling to sacrifice much for the sake of generations in the distant future.

An alternative perspective on the results comes from looking at the shape of $h(C_{Q,R})$ and $h(\hat{C}_{Q,R})$ as a function of time or decision horizon. These are shown in Figure 4. The PDFs of the climate damages, $h(C_{Q,R})$, in panels (a) and (b), look like reflected versions of $h(\Delta T)$: initially the PDFs are Gaussian, and the fat tail of worse damages grows only slowly over time. The probabilities associated with C_R are slightly more spread out than for C_Q, particularly as time progresses. On the other hand, the distributions $h(\hat{C}_{Q,R})$ converge quickly to near-steady shapes that are quite similar to each other. Both PDFs are narrowly distributed, lying predominantly in the range $99\% < \hat{C}_{Q,R} < 99.8\%$. Again, it is effective discounting of future consumption at the margin, combined with the slowly growing fat tail, that are responsible for the similarities of these PDFs.

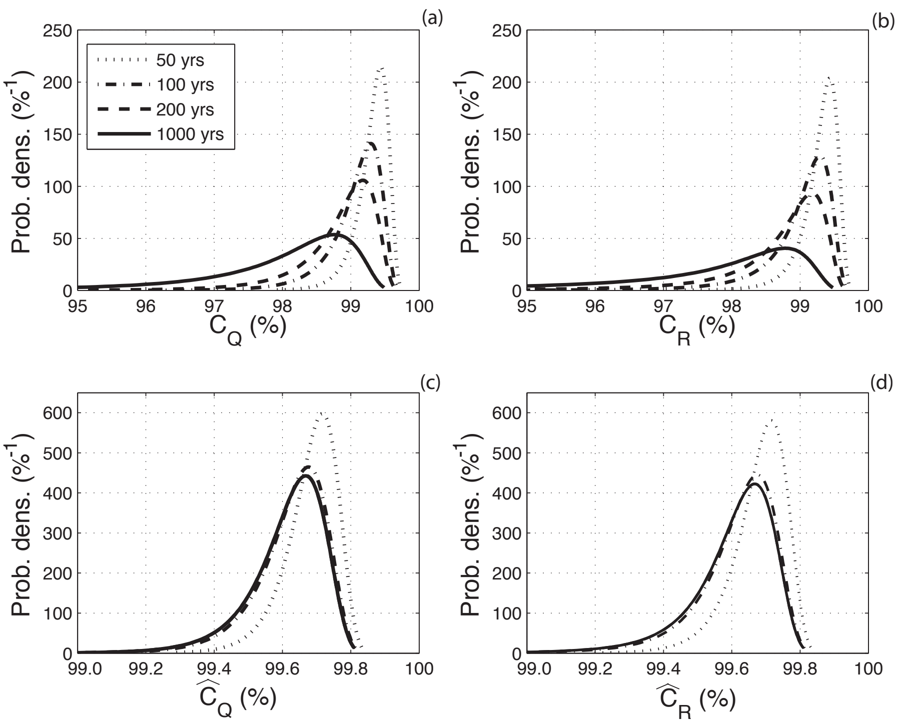

Figure 4: (a) and (b) the PDFs of climate damage function $h(C_{Q,R})$. Note that the PDFs spread out gradually over time, towards the direction of greater damages. $h(C_Q)$ is slightly more spread out than $h(C_R)$; (c) and (d) the PDFs of welfare-equivalent consumption $h(\hat{C}_{Q,R})$. Both PDFs quickly asymptote to a constant distribution because future changes are discounted in the analysis. All curves are for an instantaneous doubling of CO_2 at t =0.

For any particular forcing scenario there are two essential factors in the analysis that trade-off against each other: first is the climate sensitivity, which controls the size and growth of the fat tail;

second is the assumed long-term growth rate, which governs how strongly the future evolution of the fat tail is discounted. How does our analysis change because of uncertainty in these two factors? Figure 5 shows contours of $\hat{C}_{Q,R}(t_h = \infty)$ as a function of growth rate g, and the likelihood of a given climate sensitivity. We consider $0 \leq g \leq 3.0\%\ \mathrm{yr}^{-1}$, and consider only the far end of the fat tail of possible climate sensitivities (the upper 5%, or greater than about 2σ). Recall also that we interpret the probabilities we assign to climate sensitivity as being the maximum possible (i.e., they are an extreme upper bound).

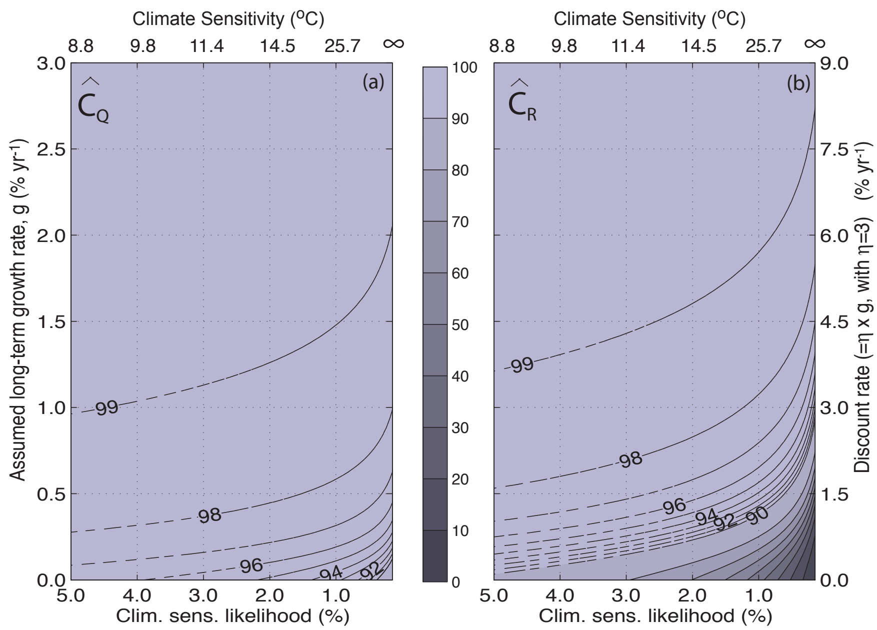

Figure 5: (a) $\hat{C}_Q(\infty)$ and (b) $\hat{C}_R(\infty)$ as a function of the assumed long term growth rate and the likelihood of the climate sensitivity. The top axis provides the values of the climate sensitivity accompanying a given likelihood (though recall these are extreme estimates). The right hand axis on panel (b) gives the discount rate ($= |\eta \times g|$). To show the structure of the solutions more clearly, two contouring styles are show. Shading denotes increments of 0.1 in $\hat{C}$, represented in the color bar. Values of $0.9 \leq \hat{C} \leq 0.99$ have line contours in increments of 0.01.

There are a couple of lessons one can draw from Figure 5. The first is that, for both $\hat{C}_Q$ and $\hat{C}_R$, it is only for low long-term growth rates (or equivalently low discount rates), and only for the extreme far tail of the climate sensitivity distribution, that the long-term damages exceed 2% of equivalent consumption. The second is that $\hat{C}_R$ diverges significantly from $\hat{C}_Q$ under the same conditions, though in a slightly smaller corner of parameter space.

4 Could it get any worse?

We've used a doubling of CO_2 as an example because of its long history as a benchmark measure of climate change, but any forcing scenario can be readily entertained. We next consider a more realistic scenario in which relatively little is done about GHG emissions and a quadrupling of CO_2 occurs over 200 years, followed by stabilization. CO_2 therefore reaches 1120 ppmv, producing a radiative forcing of 8 Wm^{-2}. If $t = 0$ corresponds to circa 1950, then this is a reasonable representation of a business-as-usual scenario out to 2150 (e.g., our simple scenario lands somewhere between the A1B and the A1F1 scenarios of AR4, two of the largest forcing projections countenanced by the IPCC). Figure 4 shows the evolution of $h(\Delta T)$, $h(C_{Q,R})$, and $h(\hat{C}_{Q,R})$. Our uncertainties in the evolution of ΔT are quite consistent with those in IPCC (2007) for equivalent scenarios.

Because climate forcing increases from zero gradually, temperature increases are actually less, initially, than for the instantaneous doubling of CO_2 (i.e., compare with Figure 1). By year 200, though, the 2σ range spans approximately 2.6 to 5.6 $^{\mathrm{o}}$C. After stabilization the mode of $h(\Delta T)$ increases only a little, but of course the fat tail at high ΔT continues to grow throughout the millennium. Except initially, the climate damages, $C_{Q,R}$, are larger for this more severe climate forcing, and after year 200 there is a radical divergence between C_Q and C_R. Even for the mode of

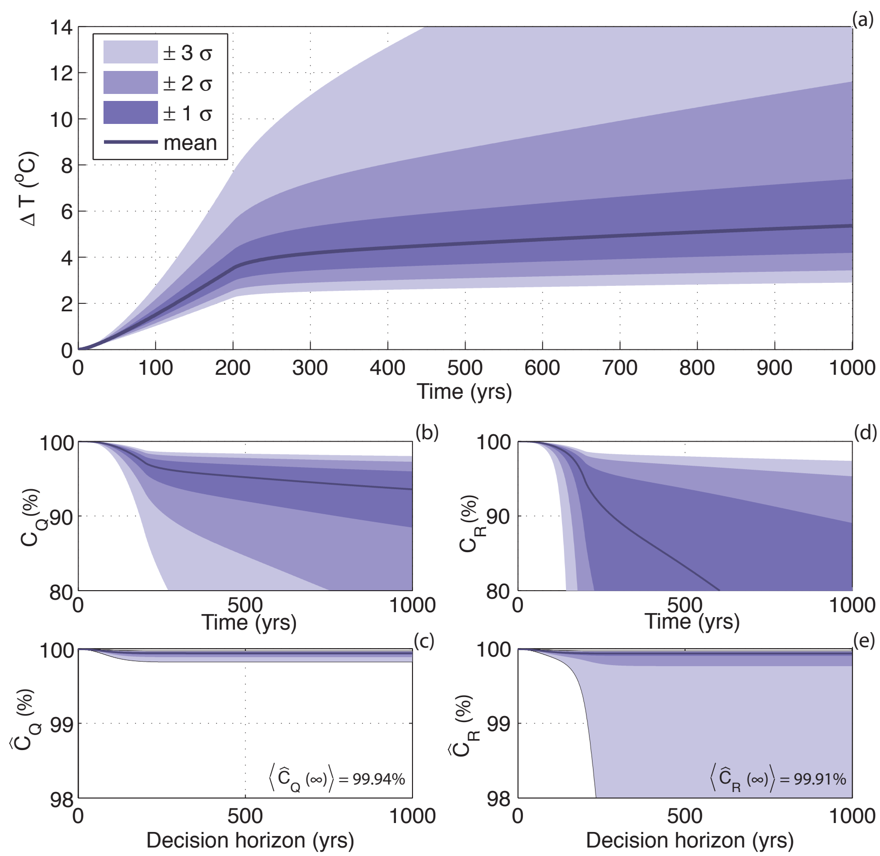

Figure 6: As for Figure 3, but for a quadrupling of CO_2 over the first 200 yrs, followed by stabilization. Note that for this more realistic scenarios the initial growth of temperatures is actually less than for an instantaneous doubling of CO_2. See text for more description.

the distribution, C_R rapidly becomes a significant fraction of welfare-equivalent consumption.

One again, however, the effective 6% yr^{-1} discounting of marginal consumption based on $g = 2\%\ yr^{-1}$ plays a very dramatic role. At the margin, consumption in earlier years is weighted so much more strongly than in later years that the smaller initial temperature increases for this scenario win out over the larger, later increases in temperature. Both $\hat{C}_Q$ and $\hat{C}_R$ are reduced by less in this scenario than in the instantaneous CO_2 doubling. And in terms of expectation values for $t_h = \infty$, both suggest climate damages of less than 0.1% of welfare-equivalent consumption.

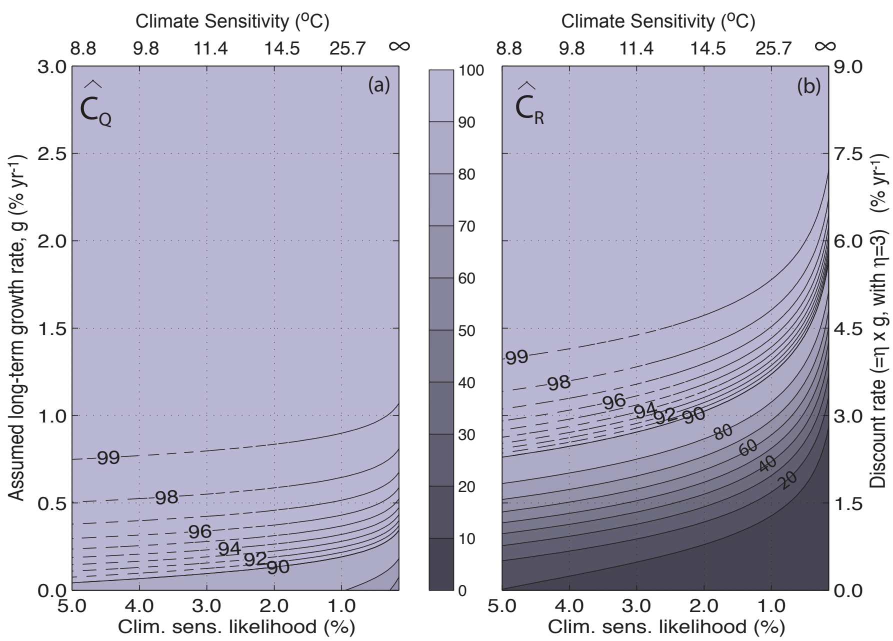

Figure 7: As for Figure 5, but for a quadrupling of CO_2 over the first 200 yrs, followed by stabilization.

Figure 7 reproduces our previous uncertainty analyses as a function of g and the likelihood of a given $\Delta T_{2\times}$ for this quadrupling scenario (i.e., compare with Figure 5). For low g ($<\sim 2.0\%\ yr^{-1}$) and for the upper reaches of $h(\Delta T_{2\times})$ tail (the upper $\sim 5\%$), there is a substantial difference between $\hat{C}_Q$ and $\hat{C}_R$. So it is for this scenario of large forcing, and this corner of parameter possibilities, that we find any real impact of the fat tail.

Finally, we summarize all of our results in Figure 8. Averaging over all possible $\Delta T_{2\times}$, we present the expectation values for $\hat{C}_{Q,R}$ as a function of g, and we also show the same calculations for the earlier, CO_2-doubling experiment. As is widely appreciated, the ramifications of these kinds of analyses depends critically on the choice for long-term growth rates. However a strong conclusion is that, if these particular choices of damage functions are to be believed, the fat tail of climate sensitivity only matters for strong climate forcing (a quadrupling of CO_2), and for low long-term growth rates ($< 1.0\%$, or equivalently, if policy makers choose not to discount the future very strongly).

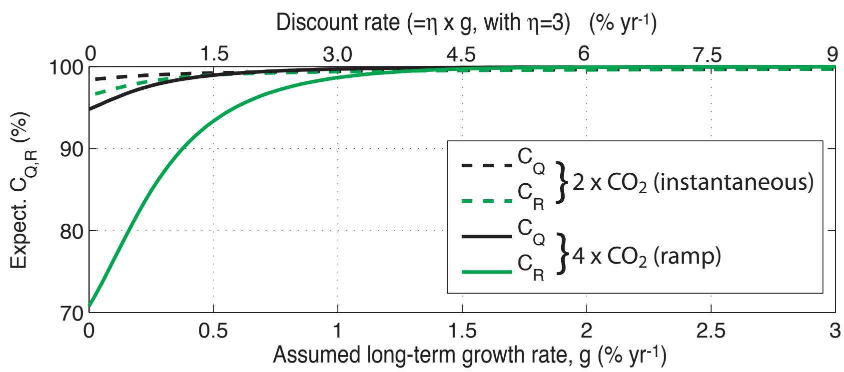

Figure 8: Summary of our analyses. The expected welfare-equivalent consumption, $\hat{C}_{R,Q}(\infty)$, as a function of assumed long-term growth rate, for the case instantaneous doubling (dashed), and ramp quadrupling (solid) of CO_2. The top axis gives the discount rate ($= |\eta \times g|$). Expectation values are calculated by averaging $\hat{C}_{R,Q}(\infty)$ over all possible $\Delta T_{2\times}$, and weighting by their relative likelihoods.

5 Discussion

Over the next several centuries, the Earth cannot reach the high temperatures that lie in the fat tail of the probability distribution of equilibrium climate sensitivity. This is true regardless of the path of anthropogenic greenhouse gases because the equilibrium is only attained when the deep ocean has warmed up. This requires millennia, and the larger the climate sensitivity, the longer it takes. In essence, this depends on little more than conservation of energy and the heat capacity of water, and is thus as certain as anything else in science. Our results spans the extreme range of possible climate sensitivity, and our conclusions are independent of the details of the shape of the distribution.

What we can say for certain is that if—and these are big "ifs"—if (1) global temperatures for the next few hundred years are driven by anthropogenic greenhouse gas emissions, and (2) significant climate impacts on human welfare occur only at large temperature increases on the order of 5°C, and (3) absent significant climate impacts, economic growth continues to raise living standards at something like $g = 2\%\ \text{yr}^{-1}$, then (4) the economic case for spending more a few percentage points of global GDP on climate mitigation is extremely weak.

Each of these four points is worth examining in more detail. Of the three "ifs", the first is the most certain. The Earth could warm up faster than our calculated range indicates only if the release of non-anthropogenic greenhouse gases, such as clathrates, is triggered (e.g., Archer, 2007). Though the conditions under which this might occur are highly uncertain, the positive feedbacks that such processes represent can be calibrated against the geologic record (e.g., Torn and Harte, 2006; Hansen et al., 2008), and readily be incorporated into simple climate models such as the one we have used. Even so, the rate of warming will still be impeded by the thermal inertia of the ocean, leaving the basic analysis here unchanged.

The second "if"—concerning the effect of temperature change on human welfare—is quite uncertain, in part because global temperature change merely serves as an ingenuous proxy for the "bad things" associated with climate change, and thus substitutes for a panoply of deleterious climate impacts. Its use is perhaps a convenience because, of all the climatic variables that one can consider, it has the closest relationship to emissions scenarios. In many respects ways, however, it is a crude tool for the task, and it may well be that focussing on a subset of environmental phenomena would enable a more precise explication of the extreme climate outcomes that are held to be unacceptable. For example, threats to freshwater resources or agriculture on a continental scale (e.g., Seager et al., 2007; Battisti and Naylor, 2009), increased frequency of injurious heat waves (e.g., Diffenbaugh and Ashfaq, 2010), ocean acidification and the health of marine ecosystems (e.g., Orr et al., 2005) are all candidates for quantifiable catastrophes that can happen on a century timescale, and for lower global mean temperatures than those residing in the fat tail of climate sensitivity. For many of these impacts the link to emissions is less certain than for global mean temperature, and therefore harder to integrate into a dynamic integrated assessment models such as DICE (Nordhaus, 2008), However, they may form a firmer basis for establishing what level of human interference in climate should be prohibited. Answers to these kind of analyses depend on the ability to make regional climate predictions, which is currently a focal topic of climate dynamics research and emphasized as an important target for the next IPCC report. It also depends on being able to define a meaningful damages function, which may be the greater challenge.

The third "if"—about long-run economic growth rates averaging $g = 2\%\ \mathrm{yr}^{-1}$—is perhaps the most uncertain. With what confidence can the consequences of our actions be exponentiated into the future? While recent growth rates may act as reasonable bounds for the next few decades, it is far from incontrovertible that per-capita consumption in 250 years will be 150 times greater than it is today, or that per-capita consumption in 500 years will be 20,000 times greater than it is today. It is worth remembering that recent growth rates are exactly that: recent. Maddison (2007) examines global economic conditions for the past 2000 years and concludes that per-capita income fell during the first 1000 years and rose by only 50% over the next 800 years, with "dynamic" growth (per-capita income rising by a factor of ten) occurring only in the period since 1820. Whether this dynamic growth rate—reminiscent of Moore's famous law for transistors—can continue is unknown, but Moore's less-well-known corollary (that "no exponential is forever") is worth keeping in mind.

Moreover, a planet that is hotter by 5 °C, or more, is a very different world and, in the opinion of many, would be a dismal legacy of economic and human progress, that would also engender a hideous disruption to other life on Earth. Powerful emotions resile against the prospect of bequeathing such a world to our descendants, but economic arguments that factor in conventional long-term growth rates are blind to such feelings. Through the lens of future generations, one can easily imagine that their increased consumption will not be the only measure by which they judge us.

Finally, there is point (4), that the economic case for spending more a few percentage points of global GDP on climate mitigation is extremely weak if you accept the first three assumptions. The most important observation here is that an economically efficient carbon-mitigation effort—a steadily rising carbon tax with revenues used to reduce existing distortionary taxes—is estimated to cost only a few percentage points of global GDP. As always, economic analysis involves comparing costs and benefits, and if the costs of emissions reduction are low then the benefits of emissions reduction (i.e., the harms of climate change) do not have to be large in order to warrant action.

6 Acknowledgements:

Marcia, Marshall

References

Allen, M.R., and coauthors, 2006: Observational constraints on climate sensitivity. In *Avoiding Dangerous Climate Change*, ed H.J. Schellnhuber, Cambridge University Press, Cambridge, U.K.

Archer, D., 2007: Methane hydrate stability and anthropogenic climate change. *Biogeosciences* **4**, 521-544.

Annan, J.D., and J.C. Hargreaves, 2006: Using multiple observationally-based constraints to estimate climate sensitivity. *Geophys. Res. Lett.*, **33**, L06704, doi:10.1029/2005GL025259 .

Annan, J.D., and J.C. Hargreavs, 2010: Reliability of the CMIP3 ensemble. *Geophys. Res. Lett.*,**37**, L02703, doi:10.1029/2009GL041994.

Arrhenius, S., 1896: On the influence of carbon acid in the air upon the temperature of the ground. *Philosoph. Mag.*, **41**, 237-276.

Baker, M.B., and G.H. Roe, 2009: The shape of things to come: why is climate change so predictable? *J. Climate*, **22**, 4574-4589.

Battisti, D.S. and R.L. Naylor, 2009: Historical Warnings of Future Food Insecurity with Unprecedented Seasonal Heat. Science, 323, no 5911, 240-4.

Charney, JG, 1979: Carbon dioxide and climate: a scientific assessment. National Academy of Science, 22 pp.

Diffenbaugh, N.S. and M. Ashfaq, 2010: Intensification of hot extremes in the United States, *Geophysical Research Letters*, L15701, doi:10.1029/2010GL043888.

Hansen, J., G. Russell, A. Lacis, I. Fung, D. Rind, and P. Stone, 1985: Climate Response Times: Dependence on climate sensitivity and ocean mixing. *Science*, **229**, 857859.

Hansen, J., Mki. Sato, P. Kharecha, D. Beerling, R. Berner, V. Masson-Delmotte, M. Pagani, M. Raymo, D.L. Royer, and J.C. Zachos, 2008: Target atmospheric CO2: Where should humanity aim? *Open Atmos. Sci. J.*, **2**, 217-231, doi:10.2174/1874282300802010217.

Knutti, R., G.C. Hegerl, 2008: The equilibrium sensitivity of the Earth's temperature to radiation changes. *Nature Geoscience*, **1**:11, 735-743.

Knutti, R., R. Furrer, C. Tebaldi, J. Cermak, G.A. Meehl, 2010: Challenges in combining projections from multiple climate models. *J. Climate*,**23**, 2739-2758.

IPCC, 2007: *Physical Science Basis.* Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge Univ. Press.

Maddison, A., 2007: *Contours of the World Economy, 1-2030 AD.* Oxford University Press, 448pp.

Nordhaus, W.D., 2008: A Question of Balance: Economic Modeling of Global Warming, Yale Press.

Orr, J.C., and co-authors, 2005: Anthropogenic ocean acidification over the twenty-first century and its impact on calcifying organisms. *Nature*, **437**(7059), 681686.

Roe, G.H. and M. B. Baker, 2007: Why is climate sensitivity so unpredictable? *Science*, **318**, 629-632.

Roe, G.H., 2009: Feedbacks, time scales, and seeing red. Annual Reviews of Earth and Planetary Sciences, 37: 93-115.

Seager, R., M. Ting, I. M. Held, Y. Kushnir, J. Lu, G. Vecchi, H.-P. Huang, N. Harnik, A. Leetmaa, N.-C. Lau, C. Li, J. Velez, and N. Naik, 2007: Model projections of an imminent transition to a more arid climate in southwestern North America, *Science*, **316**, 1181-1184.

Solomon, S. and fourteen others, 2010: Climate Stabilization Targets: Emissions, Concentrations and Impacts over Decades to Millennia. National Research Council, National Academy of Sciences.

Stainforth, D.A., T. Aina, C. Christensen, M. Collins, N. Faull, D.J. Frame, J.A. Kettleborough, S. Knight, A. Martin, J.M. Murphy, C. Piani, D. Sexton, L.A. Smith, R.A. Spicer, A. J. Thorpe, and M.R. Allen, 2005: Uncertainty in predictions of the climate response to rising levels of greenhouse gases, *Nature*, **433**, 403-406.

Torn, M.S., and J. Harte, 2006: Missing feedbacks, asymmetric uncertainties, and the underestimation of future warming, *Geophys. Res. Lett.*, **33**, L10703, doi:10.1029/2005GL025540.

Weitzman, M.L., 2009a: Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting. Economics: The Open-Access, Open-Assessment E-Journal, Vol. 3, 2009-39.

Weitzman, M.L., 2009b: On Modeling and Interpreting the Economics of Catastrophic Climate Change, *Review of Economics and Statistics*, **91**, 1-19.

Weitzman, M.L, 2010: GHG Targets as Insurance Against Catastrophic Climate Damages, *in preparation*

Wigley, T.M.L., and M.E. Schlesinger, 1985: Analytical solution for the effect of increasing CO_2 on global mean temperature. *Nature*, **315**, 649652.