Martin L. Weitzman
Notes for EPA & DOE discussion meeting
November, 2010

A. First thoughts on "'thinking about' high-temperature damages from potential catastrophes in climate change."

'Thinking about' is the right phrase. This is a notoriously intractable area even to conceptualize, much less to model or to quantify. Don't expect miracles or breakthroughs here ― too many "unknown unknowns" with seemingly non-negligible probabilities to feel comfortable with.

B. What is the nature of the beast?

The economics of climate change consists of a very long chain of tenuous inferences fraught with big uncertainties in every link: beginning with unknown base-case GHG emissions; then compounded by big uncertainties about how available policies and policy levers will transfer into actual GHG emissions; compounded by big uncertainties about how GHG flow emissions accumulate via the carbon cycle into GHG stock concentrations; compounded by big uncertainties about how and when GHG stock concentrations translate into global average temperature changes; compounded by big uncertainties about how global average temperature changes decompose into regional climate changes; compounded by big uncertainties about how adaptations to, and mitigations of, regional climate-change damages are translated into regional utility changes via a regional "damages function"; compounded by big uncertainties about how future regional utility changes are aggregated into a worldwide utility function and what should be its overall degree of risk aversion; compounded by big uncertainties about what discount rate should be used to convert everything into expected-present-discounted values. The result of this lengthy cascading of big uncertainties is a reduced form of truly enormous uncertainty about an integrated assessment problem whose structure wants badly be transparently understood and stress tested for catastrophic outcomes.

Let welfare W stand for expected present discounted utility, whose theoretical upper bound is B. Let $D \equiv B - W$ be expected present discounted disutility. Here D stands for what might be called the "diswelfare" of climate change. Unless otherwise noted, my default meaning of the term "fat tail" (or "thin tail") will concern the upper tail of the PDF of lnD, resulting from whatever combination of probabilistic temperature changes, temperature-sensitive damages, discounting, and so forth, by which this comes about. Empirically, it is not the fatness of the tail of temperature PDFs *alone* or the reactivity of the damages function to high temperatures *alone*, or any other factor *alone*, that counts, but the *combination* of all such factors. Probability of welfare-loss catastrophe declines in impact size, but key question here is: *how fast* a decline relative to size of catastrophe? When we turn to theory, it seems to highlight that the core "tail fattening" mechanism is an inherent inability to learn about extreme events from limited data.

C. What do rough calculations show about this beast?

I have played with some extremely rough numerical examples. GHG concentration implies a PDF of temperature responses implies a PDF of damages (given a "damages function"). In order to get tail fatness to matter for willingness to pay to avoid climate change requires a *much* more reactive damages function than the usual quadratic. Usual quadratic damages function loses 26% of output for a 12dC temperature change. At 2% annual growth rate, 12dC change 200 years from now implies that welfare-equivalent consumption then will still be *37 times higher than today*. If you use the standard quadratic damages function, you cannot get much damage from extreme temperatures. If make a reactive damages function, such that, say, 12dC temperature increase causes welfare-equivalent consumption to shrink to, say, 5% of today's level, then get very high WTP to reduce GHG target levels. Model is terrified of flirting with high CO2-e levels, especially above 700 ppm. Incredible dependence on degree of risk aversion (2, 3, or 4?), fatness of tail PDFs (climate sensitivity PDF: normal, lognormal, Pareto?), and so forth. My own tentative summary conclusion: tail of extreme climate change welfare-loss possibilities is much too fat for comfort when combined with reactive damages at high temperatures. It looks like this could influence such things as social cost of carbon.

D. Is there anything constructive to take away from this gloomy beast?

My tentative answer: a qualified maybe. Some possible rough ideas follow.

1. Keep a sense of balance. A small but fat-tailed probability of disastrous damages is not a realization of a disaster. Highly likely outcome is a future sense that we dodged a bullet (like Cuban missile crisis?). Yet when all is said and done, catastrophic climate change looks to me like a very serious issue.

2. Try standard CBA or IAM exercises in good faith. *But*, be prepared – when dealing with extremes – that answers might depend non-robustly upon seemingly-obscure assumptions about tail fatness, about how the extreme damages are specified (functional forms, parameter values, etc.), assumptions about rates of pure time preference, degrees of risk aversion, Bayesian learning, CO2 stock inertia, CH4 releases from clathrates, mid-course correction possibilities, etc. Some crude calculations seem to indicate great welfare sensitivity to seemingly-obscure factors such as the above, most of which are difficult to know with any degree of precision. Do CBAs and IAMs, study answers, but maybe don't try to deny the undeniable if these answers are sensitive to tail assumptions in a highly nonlinear welfare response to extreme uncertainty.

3. Should we admit to the public that climate change CBA looks more iffy and less robust than, say, CBA of SO2 abatement, or would this be self defeating?

4. Maybe there should be relatively more research emphasis on understanding extreme tail behavior of climate-change welfare disasters. Alas, this is very easy to say but very difficult to enact. How do we learn the fatness of PDF tails from limited observations or experience?

5. A need to compare how fat are tails of climate-change welfare loss with how fat are tails of any proposed solutions, such as nuclear power, below-ground carbon sequestration, etc.

6. Suppose that a lot of expected present discounted disutility is in the bad fat tail of the welfare-loss PDF. Realistically, how can we limit some of the most horrific losses in worst-case scenarios? Can we filter-learn fast enough to offset residence time of atmospheric CO2 stocks by altering GHG emission flows in time to work? Is tail fatness an argument for developing an emergency-standby backstop role for fast geoengineering? Any other backstop options? Take-home lesson here: hope for the best and prepare for the worst. At least we should be prepared, beforehand, for dealing with ugly scenarios, even if they are low-probability events. Should the discussion about emergency preparedness begin now?