Fat Tails and the Social Cost of Carbon

Martin L. Weitzman*

AER P&P: 12/07/13

At high enough greenhouse gas (GHG) concentrations, climate change might conceivably cause catastrophic damages with small but non-negligible probabilities. Other things being equal, the possibility of catastrophic damages should lower the discount rate used to evaluate mitigation-investment decisions and raise the social cost of carbon (SCC). This is not surprising. What is perhaps surprising is the potential strength of this effect. If the bad tail of climate damages is sufficiently fat with probability, and if the utility function has relative risk aversion greater than one, then (at least in theory for at least some formulations) the insurance-like catastrophe-reducing aspect of mitigation investments can be very powerful. This reflects the idea that, in such situations, a *tail-hedge investment*, which shifts the probability distribution of catastrophic damages in a good direction away from the bad tail, can be very valuable. In the most extreme limit this tail-hedge insurance effect can be infinitely strong and can dominate the economic analysis by making the SCC infinite. This kind of extreme (and empirically unbelievable) limiting result is a version of what I have previously labeled the "dismal theorem."[1]

In this paper I revisit the influence of fat tails on climate change economics via a stark two-period two-state formulation, which is focused on the SCC and lays simplistically bare the basic structure of the argument. I then attempt to place the underlying issues in a balanced perspective. The "dismal theorem" of an infinite SCC is a theoretical limiting result, which depends on particular assumptions that may or may not have actual relevance for climate-change policy depending upon the interaction of a variety of empirical factors, functional forms, and parameter values. I argue that the main value of the "dismal theorem" is to serve as a warning flag that a credible economic analysis of climate change should seriously consider extreme tail values of damages and their associated probabilities because they may have the potential to increase the SCC significantly.

*Department of Economics, Harvard University (mweitzman@harvard.edu).

[1]See Weitzman (2009).

1 A Super-Simple Expository Model of the SCC

The expository model here has two periods. Some base case of abatement strategy is given. All consumption refers to "effective consumption" – *after* climate change damages have been subtracted. The utility of consumption is $U(C)$. Present consumption is C_0. Future consumption is the random variable $\widetilde{C}$, whose expected utility is discounted by β. Welfare is $W = U(C_0) + \beta\,\mathbb{E}[U(\widetilde{C})]$, where $\mathbb{E}$ is the expectation operator.

Let A stand for an *extra* small amount of carbon abatement delivered at present time 0. For simplicity, the effect of extra abatement $A > 0$ is to uniformly shift upwards future consumption by the multiplicative function $\phi(A)$, where $\phi(0) = 1$ and $\phi'(A) > 0$. (This is essentially the same thing as having a multiplicative damages function.) Then welfare W becomes transformed into

$$W(C_0, A) = U(C_0) + \beta\,\mathbb{E}[U(\phi(A)\widetilde{C})]. \tag{1}$$

The *social cost of carbon* (SCC), is the (negative of the) change in C_0 per small change in A that would give the same level of welfare as before. In words it is the willingness to pay for an extra small unit of abatement.[2] Mathematically,

$$SCC = -\left(\frac{dC_0}{dA}\right)_W = \left(\frac{\frac{\partial W}{\partial A}}{\frac{\partial W}{\partial C_0}}\right)_{A=0}. \tag{2}$$

Now plug (1) into (2) and take the appropriate derivatives, which yields the formula

$$SCC = \frac{\beta\,\phi'(0)}{U'(C_0)}\mathbb{E}[\widetilde{C}\,U'(\widetilde{C})]. \tag{3}$$

For utmost simplicity, I now unrealistically assume a two-point distribution of $\widetilde{C}$. Suppose $\widetilde{C} = \overline{C}$ with probability $1 - p$ and $\widetilde{C} = \underline{C}$ with probability p, where p is assumed to be tiny. In the interpretation here, $\overline{C}$ represents the "expected" climate outcome while $\underline{C}$ ($<< \overline{C}$) represents a catastrophic climate outcome that might occur only with some tiny probability p. Assume utility is of the CRRA form $U(C) = C^{1-\eta}/(1-\eta)$, where the coefficient of relative risk aversion is $\eta > 1$. With this specification, the above formula (3) becomes

$$SCC = \frac{\beta\,\phi'(0)}{C_0^{-\eta}}\left[(1-p)\overline{C}^{1-\eta} + p\underline{C}^{1-\eta}\right]. \tag{4}$$

I want to analyze the extreme case of a *very* rare, *very* catastrophic climate event where

[2] What follows is an analytical abstraction of the procedure for deriving the SCC described, e.g., in Greenstone, Kopits, and Wolverton (2013).

simultaneously p and $\underline{C}$ are both small. Several observers have expressed the belief that an essential ingredient in a cost-benefit analysis of climate change is the potential for a small-probability high-impact disastrous outcome.[3] I think it is at least interesting, and may perhaps give some useful insights, to investigate what happens to the SCC of formula (4) in the most extreme limit as simultaneously $p \to 0$ and $\underline{C} \to 0$ – even if we cannot or will not go all the way to that limit in practice. It is a kind of "stress test" of the model.

To get a handle on this limiting issue and relate it to the tail fatness of a particular distribution, let us take $x = -\ln \underline{C}$ as a measure of how deep into the bad tail we are. Let $p(x)$ be the probability of x. Then $\underline{C}(x) = \exp(-x)$ and formula (4) can equivalently be rewritten as

$$SCC(x) = \frac{\beta\,\phi'(0)}{C_0^{-\eta}} \left[(1-p(x))\overline{C}^{1-\eta} + p(x)\exp((\eta-1)x)\right]. \tag{5}$$

Expressed in terms of x, we now investigate the issue of what happens to $SCC(x)$ as $x \to \infty$. The outcome depends on *how fast* $p(x) \to 0$ as $x \to \infty$. If the probability $p(x)$ declines in x *faster* than exponentially, then (abusing terminology) $p(x)$ is thin tailed and $SCC(x)$ in (5) goes to some finite limit as $x \to \infty$. This is the kind of situation that can justify a "value at risk" type calculation that would cut off the distribution for some large $\overline{x}$ (or, equivalently, small $p(\overline{x})$) and ignore what is in the bad tail for values $x > \overline{x}$. But what happens if $p(x)$ declines in x relatively slowly? Suppose (again abusing terminology) that the probability distribution $p(x)$ is fat-tailed, meaning that $p(x) \to 0$ polynomially as $x \to \infty$. Then $SCC(x)$ in formula (5) explodes as $x \to \infty$, and a "value at risk" type cutoff of the bad tail is not legitimate. I investigate this unusual and artificial situation in the next section.

2 A "Dismal Theorem"

Suppose that for large x the probability $p(x)$ is polynomial, meaning (for large x) that $p(x) \propto x^{-\alpha}$, where $\alpha > 0$. (This is the prototypical example of the relatively slow asymptotic probability convergence to zero that describes a fat tail.) Then we have the following result.

If $p(x) \propto x^{-\alpha}$ and if $\eta > 1$, then

$$\lim_{x\to\infty} SCC(x) = \infty. \tag{6}$$

I will call (6) (a form of) the "dismal theorem." Let us immediately emphasize that which is immediately obvious. *The "dismal theorem" is an absurd result!* It cannot be the case

[3] See, e.g., Barro (2013), Litterman (2013), Pindyck (2013).

that society would pay an infinite amount to abate one unit of carbon dioxide. Something must be very wrong in the formulation of the underlying model.

Several things could be seriously wrong with the underlying formulation.[4] The limiting probabilities might not be fat-tailed in the exact and demanding sense of this model. The CRRA utility function might be inapplicable, at least in the limiting range of infinitesimal consumption where it yields an unboundedly low value. The catastrophic realizations of ever-larger x might be occurring in the ever-more-distant future, so that in formula (5) one has to take a double limit as $\beta \to 0$ and $x \to \infty$. There might be something fundamentally wrong with applying the expected present discounted utility framework to such an extreme problem. I refrain from listing other possibilities. There are more than enough plausible arguments to explain away the infinity in the "dismal theorem" result (6).

Formally, it is not difficult to get rid of the infinity symbol in equation (6). One easy way is to not allow the limit in (6) to occur simply by fixing x at some finite value $x = \overline{x} < \infty$. This is analogous to what "value at risk" cutoffs of the bad tail attempt to do. But if $p(x)$ represents a fat-tailed slowly-converging polynomial distribution, then $SCC(\overline{x})$ will be sensitive to $\overline{x}$, which is not a fully comfortable resolution.

Why might $p(x)$ have a fat-tailed slowly-converging polynomial form? I do not have a good answer to this important question. We are really extremely unsure about the consequences of high levels of GHG concentrations. We have very little idea about the relevant probability distribution for the bad tail of extreme outcomes of catastrophic damages. Under such circumstances, where it is difficult to bound high-GHG damages, I think it is enough justification to simply say that we would like to be aware of the theoretical consequences of fat tails, leaving in temporary abeyance their empirical relevance.

I wish I could report decisive convincing numerical results from modeling catastrophic climate change. Alas and not surprisingly, any such results depend on the particular specifications going into a particular integrated assessment model. To get catastrophic climate change to matter for policy depends on some combination of a high-enough probability of occurrence, a high-enough level of catastrophic damages, strong-enough tail hedging, a high-enough level of risk aversion, low-enough time discounting, and several other features. Some researchers have found significant tail effects under some seemingly plausible specifications. I think the summary of Simon Dietz (2011) is fair: "To what extent does economic analysis of climate change depend on low-probability high-impact events? The short answer is a great deal, but not to the exclusion of other factors that we already know to be very important..."

[4] See Antony Millner (2013) for an enumeration and evaluation of various complaints against the "dismal theorem" that have appeared in the literature. Millner concludes with his own overall assessment, which I think is fair.

3 Conclusion: Fat Tail as Cautionary Tale

If the "dismal theorem" is a *reductio ad absurdum*, what are we left with?

Let us go back to basics. From (1), an investment in abatement is shifting upward the probability distribution of effective consumption. Instead of $\widetilde{C} = \underline{C}$ with probability p (and $\widetilde{C} = \overline{C}$ with probability $1-p$), a small unit investment in abatement makes $\widetilde{C} = (1+\phi'(0))\underline{C}$ with probability p (and $\widetilde{C} = (1+\phi'(0))\overline{C}$ with probability $1-p$), where $\phi'(0) > 0$. In other words, the decrease in damages from an abatement investment is equivalent to first-order stochastic dominance in the distribution of consumption. How much is first-order stochastic dominance in consumption worth? Potentially quite a lot if $\underline{C}$ is catastrophically low with a "fat" probability, because it pulls us away from the terrible tail with its terrible consequences. Abatement in this case represents a valuable tail-hedge insurance investment that shows itself in a high SCC. An ultra-extreme version of this tail-hedge effect drives the "dismal theorem."

Since the conditions for the "dismal theorem" to hold are unusual and open to legitimate criticisms, it cannot possibly trump all other considerations. Whether some modified version of the "dismal theorem" is relevant or not is ultimately an empirical question. Unfortunately, it is an empirical question that depends on probability assumptions about extreme tail behavior, which are very difficult to resolve because we know hardly anything about extreme tail probabilities. The nature of tail events is that we have little past experience with them, and besides, climate change is a unique one-off event. This is the fundamental "*dismal dilemma*" of fat tails in climate change. Fat tails may be very important, but how can we know their relative fatness and the tail-hedging effect of a given climate-change investment?

The "dismal theorem" and its accompanying "dismal dilemma" are best understood as a cautionary tale. A fat tail for rare disasters has the *potential* to dominate economic calculations like the SCC (even while it may be empirically difficult to determine its actual effect). Therefore, analysis of a situation that might potentially be catastrophic cannot afford to ignore tail behavior. It is not enough in such situations to look just at measures of central tendency or even just at thin-tailed probability distributions. Ignorance of the potential fatness of an extreme bad tail is not an excuse for ignoring the potential fatness of an extreme bad tail. I think this warning is the main message of the "dismal theorem."

References

[1] Barro, Robert (2013). "Environmental Protection, Rare Disasters, and Discount Rates." Mimeo, September 2013.

[2] Dietz, Simon (2011). “To what extent does economic analysis of climate change depend on low-probability high-impact events?” *Climatic Change*, 103(3), 519-541.

[3] Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton (2013). “Developing a Social Cost of Carbon for U.S. Regulatory Analysis: A Methodology and Interpretation.” *Review of Environmental Economics and Policy*, 7(1): 23-46.

[4] Litterman, Robert (2013). “What is the Right Price for Carbon Emissions?” *Regulation*, Summer 2013, 24-29.

[5] Millner, Antony (2013). “On welfare frameworks and catastrophic climate risks.” *Journal of Environmental Economics and Management*, 65 (2), 310-325.

[6] Pindyck, Robert S. (2013). “Climate Change Policy: What Do the Models Tell Us?” *Journal of Economic Literature*, 51(3), 860-872.

[7] Weitzman, Martin L. (2009). “On Modeling and Interpreting the Economics of Climate Change.” *The Review of Economics and Statistics*, 91, 1-19.

[8] Weitzman Martin L. (2011). “Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change.” *Review of Environmental Economics and Policy*, 5(2): 275-292.