

Reviewer comments on Pycroft et al., “A Tale of Tails: Uncertainty and the social cost of carbon dioxide”

This is an excellent paper that should certainly be published. It offers an innovative and well-explained approach to “tail risk” in climate uncertainty. I have just a few comments which should be addressed before publication, followed by one larger question that may be too much to handle in this paper.

Page 5: Fix Pindyck reference in note 8. Also, consider adding another sentence or two explaining the unfamiliar functional form of equation (1), either in the text or in note 9.

Page 12: I don’t understand the point made in the last paragraph above the section 2.4 subhead. (The second-to-last sentence, “the marginal change is only responsible”, particularly puzzled me.) Is it saying that drawing parameters from fatter-tailed distributions, but assuming no discontinuities, has a greater effect on the SCCO₂ than the PAGE defaults of thin-tailed (triangular) distributions plus a probabilistic but very rare discontinuity? In the same spirit as the present study, the standard PAGE model architecture could also be modified by increasing the probability of discontinuity. Some probability distribution for discontinuity in standard PAGE should produce roughly the same results as fat-tailed distributions in PAGE with no discontinuities. Which approach would be a better match to the structure of scientific information about climate risks and uncertainties?

Comparisons with existing estimates, pp. 12-13: Note that the 95th percentile result for the Pareto distribution in Table 3, \$839, is reasonably close to the 95th percentile climate sensitivity results under worst-case damage assumptions in another paper submitted to this issue, <http://www.economics-ejournal.org/economics/discussionpapers/2011-40>.

Conclusion: I don’t understand the “perhaps... more important approach” of estimating plausible weights for extreme negative events, attributed to Pindyck in note 24. How does one estimate plausible weights for those events? If this is possible, doesn’t it lead to a simpler approach to calibrating probability distributions? I thought the impossibility of such “plausible” estimates was the motivation for the discussion of fat-tailed distributions.

And finally, for the bigger question: The modeling strategy adopted here creates hybrid PDFs for the key parameters, using PAGE default assumptions (triangular distributions) for the lower half of the distribution, and grafting on three functional forms for the upper half, calibrated to specific values at the 50th and 85th or 90th percentiles. This has the possibly unintended effect of creating an intermediate zone in which the thinner-tailed distributions predict higher probability of moderately bad outcomes. In figure 1, climate sensitivities between about 3.5 and 6 are most likely to occur with the normal distribution, and least likely with Pareto (the fattest-tailed choice). The same is true with damage exponents between about 2.4 and 3.6 in Figure 2.

The Monte Carlo analysis described in this article essentially performs numerical integration over the whole range of these curves. Thus its evaluation of a fat-tailed PDF weighs greater risk far out the curve in the less likely but disastrous range, versus lesser risk in the moderately bad, higher-probability intermediate range. A different, equally defensible strategy for creating the hybrid PDFs might change this balance, potentially eliminating the perverse pattern in the

intermediate zone. For instance, would it be possible to use the same PDF up to the 85th percentile calibration point, and then fit the three alternate tails beyond that point? Without doing the calculation, it seems likely to me that this would produce a much greater spread between the results with the three curves. Would this be an equally legitimate result of your methods?

I realize that addressing this question fully would require vast new calculations and model runs, and there is more than enough here already to constitute a valuable paper. Still, it's worth pondering, perhaps for next time.

To repeat my overall judgment: the bottom line is that this is an excellent paper that advances our understanding of uncertainty and tail risk in integrated assessment modeling.