DEPARTMENT OF ECONOMICS AND FINANCE

DISCUSSION PAPER 2012-04

Cheering Up the Dismal Theorem

Ross McKitrick

MARCH 16, 2012

College of Management and Economics | Guelph Ontario | Canada | N1G 2W1
www.uoguelph.ca/economics

CHEERING UP THE DISMAL THEOREM

Ross McKitrick
Department of Economics
University of Guelph
Guelph ON N1G 2W1
ross.mckitrick@uoguelph.ca

Abstract

The Weitzman Dismal Theorem (DT) suggests agents today should be willing to pay an unbounded amount to insure against fat-tailed risks of catastrophes such as climate change. The DT has been criticized for its assumption that marginal utility (MU) goes to negative infinite faster than the rate at which the probability of catastrophe goes to zero, and for the absence of learning and optimal policy. Also, it has been pointed out that if transfers to future generations are non-infinitesimal, the insurance pricing kernel must be bounded from above, making the DT rather irrelevant in practice. Herein I present a more basic criticism of the DT having to do with its mathematical derivation. The structure of the model requires use of ln(C) as an approximate measure of the change in consumption in order to introduce an e^x term and thereby put the pricing kernel into the form of a moment generating function. But ln(C) is an inaccurate approximation in the model's own context. Use of the exact measure completely changes the pricing model such that the resulting insurance contract is plausibly small, and cannot be unbounded regardless of the distribution of the assumed climate sensitivity.

JEL Codes: Q2, Q3, Q4
Key Words: climate change; insurance; catastrophe

Cheering Up the Dismal Theorem

1 Introduction

Weitzman (2009) analyzed climate change as an insurance problem, in which agents can decide to transfer part of today's income to the future, to insure against a loss of consumption. If a non-zero probability of catastrophic climate change is allowed and marginal utility goes to infinity as consumption goes to zero, the so-called "Dismal Theorem" (DT) seems to suggest that agents today would be willing to pay an unbounded amount to ensure a positive consumption level in the future. This suggests that the risk of catastrophic change should imply a much larger optimal allocation of current income towards climate policy than is indicated in conventional Integrated Assessment Modeling (e.g. Nordhaus 2009). While growth models had already shown that the combination of uncertainty and $-\infty$ marginal utility at the zero bound on consumption leads to optimal 'over-saving' compared to the certainty case (Nyarko and Olson 1996), the DT seems to show that these conditions make ordinary cost-benefit analysis and, by implication, optimal planning, altogether impossible.

This result has been criticized in some recent published (Nordhaus 2011), and as yet unpublished (Horowitz and Lange 2009, Karp 2009, Millner 2011) working papers. Nordhaus (2011) reviews the debates over how fat are the tails of the distribution of catastrophe probabilities, and points out that a lot of restrictive conditions must be met for the DT to be relevant, including an absence of learning.

Horowitz and Lange (2009) argued that as long as there exists a safe investment that generates a non-zero return, the stochastic discount factor in the Weitzman model cannot be unbounded. They also note that the Weitzman model values an infinitesimal transfer from the present to a future world where consumption would otherwise be zero at infinity. Focusing on this particular feature of his model then leads to a debate over whether the rate at which marginal utility goes to infinity swamps the rate at which the probability of the zero consumption outcome vanishes. Horowitz and Lange (2009) and Karp (2009) both find this an uninteresting question, since for a non-infinitesimal transfer, the marginal utility evaluated at the post-transfer level must be finite. Hence, as long as cost-benefit analysis is applied to non-infinitesimal policies, the DT is irrelevant.

Millner (2011) provides a detailed survey of these and other critiques of the DT. He argues that the concerns of Horowitz and Lange (2009) and Karp (2009) can be met by extensions to the original Weitzman framework, such as making the success of the transfer a random variable governed by the amount of climate change, or by picking a large enough damage function parameter. Millner raises a potential problem of Weitzman's Bayesian framework, namely the assumption of an uninformative Jeffrey's prior. Since the uninformative prior is represented by a distribution that is itself fat-tailed, the derivation of a fat-tailed posterior may therefore be tautological. Millner, however, finds that while the Jeffrey's prior is reasonable for the case at hand, so would be a maximum entropy prior which would yield a thin-tailed posterior. Millner also argues that consideration of simple representations of climatic processes indicates that even if the climate sensitivity parameter is unbounded, the observed temperature changes must themselves be bounded at any finite time. On the other hand, he argues that appropriate choices of damage function parameters may restore DT-like conclusions, namely that a positive probability exists that consumption may be driven to zero under plausible scenarios. Finally, Millner points to inherent flaws in the use of a constant relative risk aversion (CRRA) utility function, especially that the DT set-up pushes the analysis outside the limits for which they are analytically useful. He suggests consideration of alternative utility functions, and also adoption of a social choice perspective that might yield valuation of catastrophes as infinitely bad even if individual valuations are finite.

My analysis herein is concerned with a more fundamental, and indeed simple, problem in the DT model. Weitzman represents today's consumption as $C_0 \equiv 1$, and future consumption net of damages due to climate change as C. Then C/C_0 is one plus the percent change in consumption, and he uses $\ln(C)$ to approximate the percentage change. Since $e^x \approx 1 + x$ for small x this is a valid approximation for small changes in C. But it is not valid for changes larger than about ±15% (see Figure 1), namely for the kinds of larger changes on which the DT analysis is focused. The exact measure of consumption change is $(C/C_0 - 1)$. But as I show in the next section, use of this term eliminates all the basic results in the DT set-up. It leads to the conclusion is that the optimal insurance contract is much smaller, and much more plausible, regardless of the distribution of the climate sensitivity parameter. Moreover it is not possible for the price of insurance to be unbounded unless we know both that future consumption will be zero and there is no possibility of transferring wealth to the future, in which case there would be no reason to purchase insurance anyway.

DT critics like Millner (2011) have already shown that plausible alternative assumptions about specific features of the Weitzman model can yield different, and non-dismal, outcomes. My point herein is that, even staying in Weitzman's framework but using an exact measure of consumption growth rather than an approximation, unravels the model's conclusions.

2 Use of an Exact Measure of Consumption Change

In the model, today's consumption is denoted $C_0 \equiv 1$, and future consumption net of damages due to climate change is C. Utility is of the CRRA form $U = C^{1-\eta}/(1-\eta)$. The percentage change in consumption is denoted Y and is assumed to be a linear function of the temperature change, though that specific aspect of the model is not important; all that matters is that Y is a random variable. Weitzman (2009) approximates Y using $\ln(C)$, or rather approximates C using $\exp(Y)$. Suppose we want to transfer an amount g from the present to the future where it will be worth h. Denoting the discount factor as β and using $\boldsymbol{E}$ to denote expectation, a utility-neutral transfer is $U(1-g) + \beta\boldsymbol{E}U(C+h) = U(1) + \beta\boldsymbol{E}U(C)$. The limit as $g \to 0$ yields $g = (\beta\boldsymbol{E}U'(C)/U'(1))h$ (Millner 2011). So the marginal rate of substitution, or "stochastic pricing kernel" is written

$$M(C) = \beta\frac{U'(C)}{U'(1)} = \beta\exp(-\eta Y)\,. \tag{1}$$

If $Y \sim f(y)$ where y is a realization of Y then the amount of present consumption one would give up to secure an additional sure unit of future consumption is given by the expectation

$$E(M) = \beta\int_{-\infty}^{\infty} e^{-\eta y} f(y)dy \tag{2}$$

which is also the moment generating function (MGF) for f. Weitzman's formulation then proceeds as follows. Suppose s represents the variability of Y (climate sensitivity, or alternatively the standard deviation of Y), and $(y-\mu)/s \sim \phi(0,1)$ for known μ. ϕ can be any piecewise-continuous PDF including the uniform or the normal, subject to minimal regularity conditions. If a sample of evidence about y is obtained from climatic studies, the uninformative prior $p_0(s) \propto s^{-k}$ yields a posterior distribution *f* in the form of a Student's *t*, which implies that (2) is undefined, or $+\infty$. The unbounded willingness to pay to insure future consumption implies that any non-extreme policy response to, in this case, the risk of climate change, must implicitly or explicitly depend on an arbitrary truncation of the range of risks.

The use of $C = \exp(Y)$ to approximate consumption change has the effect of getting the exp(.) term into the integral in (2), which yields a result in the form of an MGF. But it is not a necessary step, and indeed it is a bit of a contrivance. It is inaccurate in the cases of interest, namely where *Y* is large, yet its removal substantially changes the model.

Y is defined as the percent change in consumption, and current consumption is unity, so

$$Y = C - 1 \Leftrightarrow C = 1 + Y \,. \qquad (3)$$

$C = \exp(Y)$ is only an approximation for small *Y* whereas (3) is exact in all cases. (1) then becomes

$$M(C) = \beta(1+Y)^{-\eta} \qquad (4)$$

so (2) becomes

$$E(M) = \beta\int_{-\infty}^{\infty} (1+y)^{-\eta} f(y)dy \qquad (5)$$

which is no longer an MGF. Integrating by parts yields

$$\frac{E(M)}{\beta} = \frac{f(y)(1+y)^{(1-\eta)}}{1-\eta} - \int \frac{(1+y)^{1-\eta}}{1-\eta} f'(y)dy \qquad (6).$$

The set-up in Weitzman (2009) treats Y as a linear function of a variable Z that follows ϕ, which, as noted, can be Uniform. In fact the Uniform case is well-suited to the consideration of worst-case climate scenarios. Rather than applying Bayes' theorem to an uninformative prior, we can simply note that the worst conceivable *posterior* range of bad outcomes (in terms of the fattest possible tails) over Y would follow $u(-1,0)$, namely losses ranging uniformly from zero to 100% of consumption. We could consider wider ranges but it would not have any effect on the results herein. As long as f is uniform over any range, $f' = 0$ and $f(y)=1$ over the support of y, hence (6) reduces to

$$E(M) = \beta C^{1-\eta} / (1-\eta) = \beta U(C) \quad (7).$$

Denoting current utility as U_0, the price of a sure unit of future consumption as a fraction of current utility can be interpreted as the fraction of current real income we would give up to fully insure future consumption, which equals

$$\frac{E(M)}{U_0} = \beta C^{1-\eta} . \quad (8)$$

If we take $\eta = 2$ (as per the illustrative cases in Weitzman 2009 and Nordhaus 2009) the premium cost is β / C. Only in the case where C=0 would this go to $+\infty$.

The difference is that in the original Weitzman result, the fact that the distribution of possible outcomes includes C=0 makes the price of insurance go to infinite, even if C=0 is only a low-probability event. (8) implies that the willingness to pay for insurance would be unbounded only if it were known with certainty that C=0 would be realized. But any insurance purchase today implies C>0 in the future, so the pricing kernel must be finite.

3 Discussion and Conclusion

The difference between (7) and (2) arises from the use of ln(C) as an approximation to the percent change in consumption. In an application where large changes in C are considered, it is not an accurate approximation, and since the DT derives its interest from the possibility of a large change, use of (3) is preferred. But it also implies that the willingness to pay even for large potential changes in future consumption must be small.

Weitzman (2009) defines the "future" as about two centuries ahead, so the discount factor β will be very small indeed. A 2.3% discount rate over 200 years yields β=0.01. If we use $\eta = 2$ we can graph $E(M)/U_0$ against the future value of C according to (8), which is shown in Figure 1. A 50% loss of income implies $E(M)/U_0$=0.02, in other words we would be willing to spend 2% of real income today to insure against a 50% loss of income 200 years ahead. For expected losses up to about 25% of income, the insurance premium is below 1.4% of income. Use of a 3.5% discount rate means these insurance premiums should be multiplied by 0.1, and so forth.

References.

Horowitz, John and Adreas Lange (2009) "What's wrong with infinity? A note on Weitzman's dismal theorem." Unpublished manuscript http://faculty.arec.umd.edu/jhorowitz/weitzman_final.pdf.

Karp, Larry (2009) "Sacrifice, discounting, and climate policy: five questions." Unpublished manuscript http://are.berkeley.edu/~karp/FiveQuestionsSubmit.pdf.

Millner, Antony (2011) "On welfare frameworks and catastrophic climate risks." Unpublished manuscript.

Nordhaus, William D. (2009) "An Analysis of the Dismal Theorem" Cowles Foundation Discussion Paper No. 1686.

Nordhaus, William D. (2011) "The Economics of Tail Events with an Application to Climate Change." *Review of Environmental Economics and Policy* volume 5, issue 2, summer 2011, pp. 240–257 doi:10.1093/reep/rer004

Nyarko, Yaw and Lars J. Olson (1996) "Optimal Growth with Unobservable Resources and Learning." *Journal of Economic Behavior and Organization* (29) 465—491.

Weitzman, M. (2009) "On Modeling and Interpreting the Economics of Catastrophic Climate Change" *Review of Economics and Statistics* 91(1) 1—19.

Figures

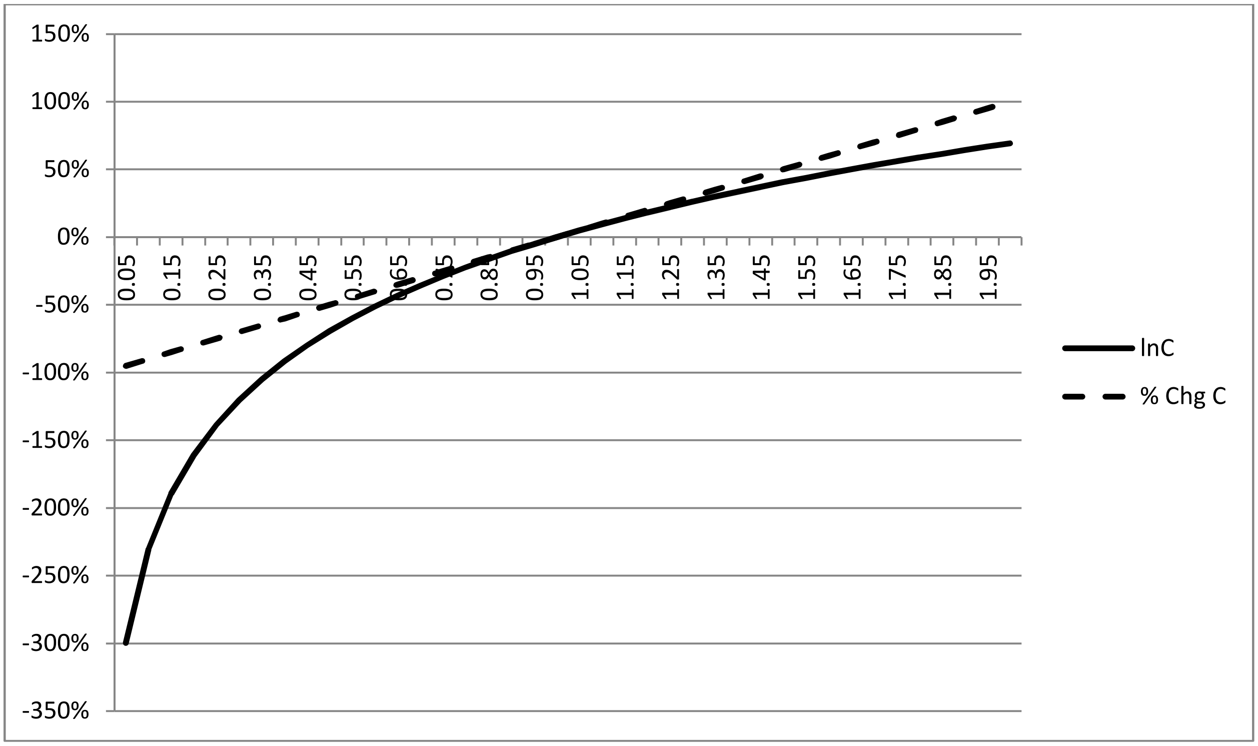

Figure 1. Horizontal axis: $\frac{C_1}{C_0} - 1$. Dashed line: % change. Solid line: ln$C$ approximation.

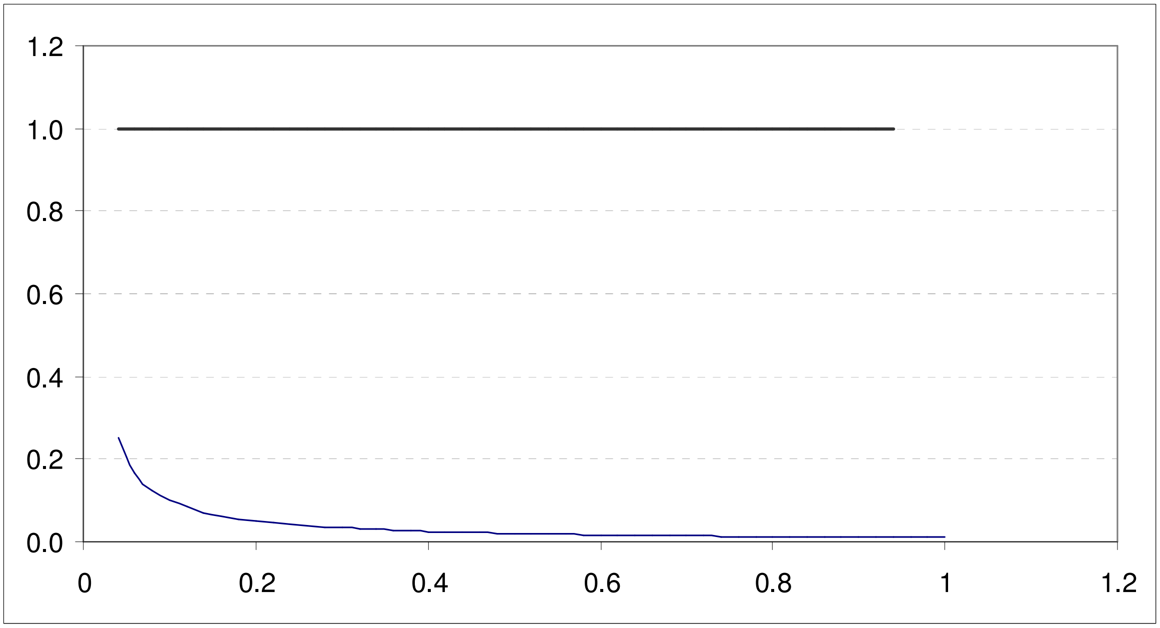

Figure 2. Willingness to pay for full insurance as a fraction of current real income ($E(M)/U_0$, vertical axis) graphed against fraction remaining of future income (C, horizontal axis) after damages. Current consumption (1.0) is shown as solid horizontal line. Assumed discount rate is 2.3%.