CHEERING UP THE DISMAL THEOREM

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

Abstract

Weitzman's (2009) Dismal Theorem (DT) states that peoples' willingness to pay to insure against possible future global warming damages is, effectively, infinitely large, as long as they have constant relative risk aversion and an uninformed prior view about the risks of climate change. This unusual result arises due to a mathematical curiosity in stochastic finance, in which the pricing rule in seemingly ordinary insurance problem takes the form of the moment generating function of a *t* distribution, which is undefined, leading to a degenerate solution. The DT has been used to argue that conventional cost-benefit analysis won't work in the climate change case, or that much more costly policy interventions than are typically considered should be pursued. I show that the degeneracy is not a general result. It relies on the use of a log approximation to the rate of consumption growth and an extreme prohibition on intergenerational wealth transfers. Use of an exact growth measure simplifies the pricing model such that insurance prices can only be unbounded in a trivial case. In general the model implies unexceptional willingness to pay to avoid future costs, even when climate risks are large and uniformly distributed.

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

Cheering Up the Dismal Theorem

1 Introduction

Weitzman (2009) analyzed climate policy as an insurance problem, in which agents can purchase a contract today to endow a future generation with full compensation in the event of warming-related damages. This is not the typical framework for studying the problem: usually it is framed in terms of assessing the potential future benefits of emission reductions against their discounted current costs. Because of the long time scales, deep uncertainties and inelastic energy demands, economists have tended to argue that only modest emission cuts can be justified on cost-benefit grounds, at least to begin with (e.g. Nordhaus 2008). Weitzman, by contrast, came to the surprising conclusion that very large costs should be incurred today to insure against future consumption losses. If a non-zero probability of catastrophic climate change exists and marginal utility goes to infinity as consumption goes to zero, the so-called "Dismal Theorem" (DT) says that, in the absence of prior information that rules out the upper tail of climate risks, agents today would be willing to pay an infinite amount to ensure a positive consumption level in the future. This suggests that the risk of catastrophic change, however remote, demands 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 infinite marginal utility at the zero bound on consumption leads to optimal 'over-saving' compared to the certainty case (Nyarko and Olson 1996), the DT seemed to show that these conditions make ordinary cost-benefit analysis and, by implication, optimal planning, altogether impossible for climate policy.

The DT has been criticized in some recent published (Nordhaus 2011, Millner 2013), and as yet unpublished (Horowitz and Lange 2009, Karp 2009) papers. These authors all accept the derivation as given but dispute the relevance of the conclusions. Nordhaus (2011) reviews debates over how fat the tails of the distribution of catastrophe probabilities are, and points out that some 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 result hinges on whether the rate at which marginal utility goes to infinity as consumption goes to zero exceeds 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, not least because it concerns small differences in slopes of ill-defined functions in unobservable regions of degenerate points, but also because as long as non-infinitesimal transfer of wealth is possible, the marginal utility evaluated at the post-transfer level must be finite. Hence, as long as cost-benefit analysis is applied to non-infinitesimally small policies, the DT is irrelevant.

Millner (2013) 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. In other words, if the world is allowed to be sufficiently dismal, the Dismal Theorem may still apply. Millner instead raises a potential problem of Weitzman's probability framework. A Bayesian model supposes that agents approach an uncertain situation with prior beliefs that are modified in light of available data to yield a posterior distribution of possible outcomes. A common way of starting such an analysis is to assume a so-called Jeffrey's prior, in which the starting knowledge is uninformative. But then the derivation of a fat-tailed posterior distribution function may be tautological. Millner argues that a common alternative formulation, called a maximum entropy prior, would also be reasonable but would yield a thin-tailed posterior distribution, thus sufficiently constraining the risks of extreme outcomes to eliminate the DT. Millner also argues that even if the climate sensitivity parameter is unbounded, the observed temperature changes must themselves be bounded at any finite time. In that case, special choices of damage function parameters are needed to generate a positive probability 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.

At this point the debate begins to sound understandably esoteric, yet the underlying concept of the DT is important and sufficiently striking that it merits thinking through. The analysis herein goes in a different direction, putting the discussion onto what is intended to be a simpler and more transparent footing so as to clarify why the DT should properly be seen as an extreme, and implausible, special case of a model that generally yields ordinary conclusions. The starting point os this analysis is an innocuous change in a preliminary element of 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 large changes on which the DT analysis is focused. An exact measure of the percentage change in consumption is $(C/C_0 - 1)$. I will show that use of this measure fundamentally changes the basic structure in the DT set-up. The price of insurance only becomes unbounded in a special case in which a market for insurance is impossible anyway, namely if we know that future consumption will be zero and there is no possibility of transferring wealth between periods, and even then only for some values of the risk aversion parameter. Thus, staying in Weitzman's utility framework but using an exact measure of consumption growth rather than an approximation, the dismal theorem only emerges as a special, and implausible, case, whereas much less dismal outcomes are the norm.

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 climatic 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 E to denote expectation, a utility-neutral transfer is

$$U(1-g)+\beta EU(C+h)=U(1)+\beta EU(C).$$

The limit as $g \to 0$ yields $g = (\beta EU'(C)/U'(1))h$ (Millner 2013). So the marginal rate of substitution, or "stochastic pricing kernel" is written

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

If the probability density function of Y is $f(y)$, 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

$$E(M) = \beta \int e^{-\eta y} f(y) dy \qquad (2).$$

This function coincides with what is known in mathematical statistics as 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 $\frac{y-\mu}{s} \sim \phi(0,1)$ 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. But the MGF of the t distribution is undefined, which implies that (2) can be regarded as infinite for any CRRA parameter. The unbounded willingness to pay to insure future consumption follows from the uninformative prior and the uncertainty over y. It 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 various attempts to get around this proceeded as described above. Here I revisit the the use of $C = \exp(Y)$ to approximate the consumption change. This 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 inaccurate in the cases of interest, namely where Y is large. Its removal, however, substantially changes the model.

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

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

Equation (3) is exact in all cases. Using it in (1) yields

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

so (2) becomes

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

where the bounds (a,b) refer to the support of f; in other words the range of the percent change in future consumption due to the risks of climate change. This expression is no longer an MGF so Weitzman's basic result, that $E(M)$ is infinitely large under general conditions, no longer follows. Integrating (5) by parts yields

$$\frac{E(M)}{\beta} = f(y)U(1 + y)|_a^b - \int_a^b U(1 + y)f'(y)dy \qquad (6).$$

The set-up in Weitzman (2009) treats Y as a linear function of a variable Z that follows the distribution function ϕ, which, as noted, can be Uniform. In fact the Uniform case is well-suited to the consideration of worst-case climate scenarios, since it implies that the only knowledge we have about the future effect of climate change is the range of potential changes, which can be as wide as we choose to consider.

If f is the uniform probability density function, $f' = 0$ and $f = (b - a)^{-1}$, so equation (6) becomes

$$E(M) = \frac{\beta}{1-\eta} \frac{(1+b)^{1-\eta} - (1+a)^{1-\eta}}{b-a} \qquad (7).$$

Any number of cases can be examined by varying the values of η, a and b. We rule out the case of b=a since that contradicts the assumption that y is uncertain. If we set $\eta = \frac{1}{2}$ then

$$E(M) = 2\beta \frac{\sqrt{(1+b)}-\sqrt{(1+a)}}{b-a} \; .$$

In this case there is no possibility of an unbounded solution for $E(M)$, even within the CCRA framework, with or without transfers to the future and irrespective of the damage function. For instance, in the catastrophic case, where a=-1 and b=0, implying expected change in consumption ranging uniformly from 0 to -100%, we obtain $E(M) = 2\beta$.

3 Deriving an Infinite Insurance Premium

In order to get an unbounded insurance premium, we need to assume $\eta > 1$ and a=-1 so that 1+a=0 becomes a denominator in equation (7). The first assumption is innocuous. If we assume $\eta = 2$, (as per the illustrative cases in Weitzman 2009 and Nordhaus 2009) then equation (7) becomes

$$E(M) = \beta\left(\frac{1}{1+a} - \frac{1}{1+b}\right)\left(\frac{1}{b-a}\right) \qquad (8).$$

In this case a=-1 implies $E(M) = \infty$. But, as pointed out in the introduction, this is s extreme as to be uninteresting. As long as a finite amount can be transferred to the future, it is possible to guarantee 1+a>0, and if nothing can be transferred to the future then insurance is impossible anyway, in which case the whole model ceases to make sense. If the problem under examination were, say, an imminent asteroid collision that was guaranteed to extinguish all life on Earth, it would be foolish to argue that since the marginal utility of consumption is infinite at the zero bound we should invest all our wealth in an insurance contract, naming as beneficiaries the generation expected to be on Earth after the asteroid hits, since they will all be dead, along with the insurance adjusters! The model is therefore discontinuous when a=-1 since the institutions on which the insurance market depends vanish. The willingness to pay for the insurance policy is therefore zero, not infinite. An unbounded insurance premium only arises in a context in which there is no demand for insurance.

We do not lose any of the "catastrophic" aspect of the underlying climate change story by imposing a bound on the consumption loss to ensure $1+a>0$, so that equation (8) remains tractable. We can then ask what we would be willing to pay to insure against consumption losses ranging from, say, 25% to 95% of current consumption: surely as catastrophic a range of possibilities as anyone could ever imagine being attributable to global warming. Weitzman (2009) defines the "future" to be about two centuries ahead. A 2.3% discount rate over 200 years yields $\beta = 0.01$. Suppose we cap the upper bound on the growth of consumption over the next 200 years at 50% (b = 0.5). Then equation (8) implies the willingness to pay to insure against a potential 25% lower-bound loss of future income is equal to 0.9% of current consumption. If the lower bound drops to a 95% loss the willingness to pay rises to 13.3% of current income (see Figure 2). By any reckoning, a potential 95% loss of consumption for the entire planet is catastrophic, and consistent with even the most pessimistic damage function parameters, yet does not imply a degenerate solution of the insurance pricing problem.

4 Discussion and Conclusion

The difference between equations (7) and (2) arises from the use of $\ln(C)$ as an approximation to the percent change in consumption. It is inaccurate in applications where large changes in C are considered, and since the DT derives its interest from the possibility of large consumption changes, use of the exact expression in equation (3) is a better choice. In that case, under a uniform distribution of possible consumption losses due to global warming, the willingness to pay to insure even against large potential future losses turn out not to be unduly large. An unbounded willingness to pay no longer emerges as a general result in the model, and in fact only emerges in the extreme case in which the insurance market itself disappears, implying zero willingness to pay.

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 (2013) "On welfare frameworks and catastrophic climate risks." *Journal of Environmental Economics and Management* Volume 65, Issue 2, pp 310–325.

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

Nordhaus, William D. (2008) *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven: Yale University Press.

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

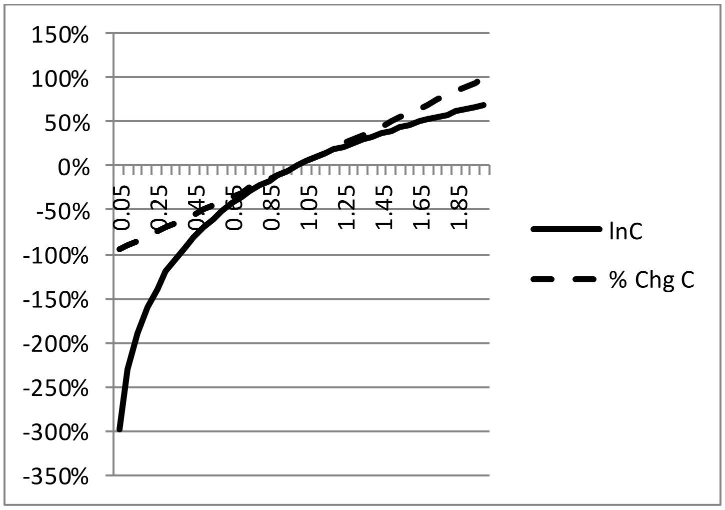

Figure 1. Horizontal axis: $\frac{C_1}{C_0} - 1$. Dashed line: % change. Solid line: $\ln C$ approximation.

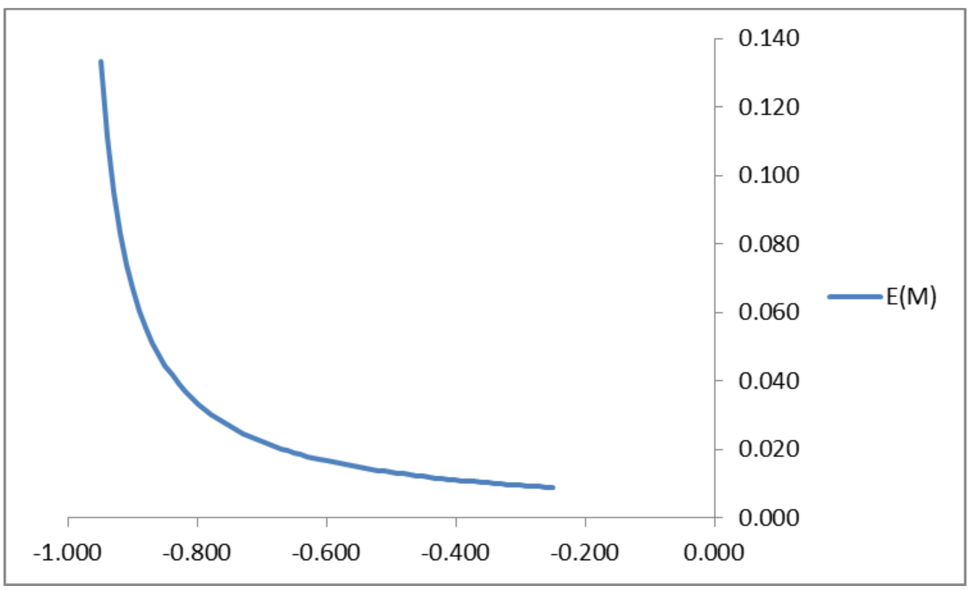

Figure 2. Willingness to pay to fully insure future income as a fraction of current real income ($E(M)$, vertical axis) graphed against lower bound on fraction loss of future income after damages over range (-0.95, -0.25). The assumed discount rate is 2.3% over 100 years and the CRRA parameter η is 2.