7

Precaution and a Dismal Theorem: Implications for Climate Policy and Climate Research

Gary W. Yohe and Richard S. J. Tol

7.1 INTRODUCTION

Economic efficiency has long been a gold standard for evaluating policies. In the context of climate change, the search for efficient solutions to the policy problem began in earnest with Nordhaus (1991), and it has evolved into using elaborate, regionally disaggregated integrated assessment models to judge the relative expected benefits and costs of various policy options across a wide range of possible futures. Cline (1992, 1997, 2004), Maddison (1995), Nordhaus (1991, 1993, 1994), Nordhaus and Yang (1996), Nordhaus and Boyer (2000), Roughgarden and Schneider (1999), Stern *et al*. (2006), Tol (2002) and Uzawa (2003) are all examples of this approach. These and many other studies are fundamentally optimization exercises, and many use Monte Carlo simulations to set the expected marginal benefits of emission reduction equal to its expected marginal cost. This is why calculations of the social cost of carbon (SCC) have become so popular.[1]

It is widely known that published estimates of the social cost of carbon vary widely. An early survey conducted by Tol (2005) reported that fully 12 % of then available published estimates were non-positive. Their median was $13 per tonne of carbon, and their mean was $85 per tonne. Tol (2007) offers an updated survey of more than 200 estimates. His new results show a median for peer-reviewed estimates with a 3 % pure rate of time preference and without equity weights of $20 per tonne of carbon with a mean of $23 per tonne of carbon. Moreover, he reports a 1 % probability that the social cost of carbon could be higher than $78 per tonne given the same assumptions, and he notes that the estimates increase rapidly as the assumed discount rate falls. Tol (2007) thereby suggests at least one reason why the range of estimates of the social cost of carbon is so large.

Gary W. Yohe: Woodhouse/Sysco Professor of Economics, Wesleyan University, Middletown, CT, USA. Contact: Department of Economics, 238 Church Street, Middletown, CT, USA 06459; gyohe@wesleyan.edu.
Richard S. J. Tol: Economic and Social Research Institute, Dublin, Ireland; Institute for Environmental Studies, Vrije Universiteit, Amsterdam, The Netherlands; Department of Spatial Economics, Vrije Universiteit, Amsterdam, The Netherlands; Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA. Contact: ESRI, Whitaker Square, Sir John Rogerson's Quay, Dublin 2, Ireland; richard.tol@esri.ie.

[1] The social cost of carbon is defined as the (expected) discounted marginal damage of carbon emissions at any point in time. Estimates of the SCC can, therefore, be interpreted as estimates of the (expected) marginal benefit of mitigation.

Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy Edited by Hélyette Geman

Hope (2006) provided some additional insight derived from exercising his PAGE2002 model. He reported that the choice of discount rate and the incorporation of equity weights are extremely important, and both lie within the purview of decision-makers. High discount rates sustain low estimates because future damages become insignificant. Conversely, low discount rates produce high estimates because future damages are important. Meanwhile, strong equity weighting across the globe support high estimates because poor developing countries are most vulnerable. Alternatively, weak or no equity weighting can produce low estimates because poor developing countries do not factor heavily in the overall calculation. Hope (2006) concluded, however, that the climate sensitivity (i.e. the increase in global mean temperature that would result from a doubling of greenhouse gas concentrations from pre-industrial levels) is the largest source of variation. It is possible to derive high estimates for the social cost of carbon even if with low discount rates and/or almost no equity weighting. All that is required is the assumption that the climate sensitivity lies at the high range of the latest range of estimates.[2]

For present purposes, it is enough to recognize that the range of estimates of the SCC of carbon is enormous for a variety of reasons – some related to decisions that human beings make in their decision process, and some related to decisions over which "Mother Nature" has purview. As a result, the cost-benefit approach to climate policy has long been vulnerable to concerns about its ability to handle adequately the scope of the underlying uncertainties and diversities of opinion.

Results drawn from the optimization approach have also been suspect because many of the potential impacts of climate change (particularly non-market impacts and low-probability but high consequence ramifications of abrupt climate change) cannot easily be quantified in economic terms. The basis of this critique of incomplete and perhaps infeasible coverage is best visualized in a matrix presented by Downing and Watkiss (2003) that tracks the degree to which the complication of climate change science is captured by benefit analysis. Three rows catalog coverage of scientific uncertainty from relatively well-established (although still uncertain) trends in climate change (e.g., average temperature, sea level rise) through considerations of the bounded risks of extreme events (including precipitation events on both sides of the distribution) and other manifestations of climate variability, and finally into representations of possible abrupt change and/or abrupt impacts. Three columns catalog coverage of economic uncertainty from relatively well-established coverage of market impacts through less robust economic assessments of non-market impacts, and into socially contingent impacts (e.g. abrupt social, political or economic changes driven by famine, migration across national borders, etc.) across multiple metrics that cannot always be quantified in economic terms. Yohe and Tirpak (2008) report that the economic analyses required to inform fully the cost-benefit approach to global climate policy has adequately covered very few of the nine combinations and permutations in the matrix.

It must be emphasized, however, that neither of these sources of concern about the applicability of the cost-benefit apparatus to climate policy is really new. Indeed, both have long histories in the literature. Early on, authors like Alcamo and Kreileman (1996), Toth *et al*. (1997) and Swart *et al*. (1998) responded to them by arguing in favor of taking a precautionary approach to climate policy – i.e., defining the boundaries of "tolerable" climate impacts calibrated in terms of temperature targets (both absolute levels and sometimes

[2] One might, for example, take climate sensitivity to be greater than 5°C and only be at the 80th percentile of the distribution reported in Andronova and Schlesinger (2001).

rates of change) and working from there. In this context, policy designers ask economists simply to calculate emissions (reduction) paths that would avoid the proscribed boundaries of climate change at minimum expected cost.[3] Wigley *et al*. (1996) and Manne and Richels (1997) are perfect examples of this type of analysis.

Many of the remaining issues for the precautionary approach pertain to defining the boundaries of tolerable climate change (or, in the parlance of the United Nations Framework Convention on Climate Change, the boundaries of "dangerous anthropogenic interference with the climate system") and coping with adaptation; see, for example, Yohe and Toth (2000). It must be noted as well, however, that the precautionary approach is not immune from its own vulnerability to enormous uncertainty. Both Stern *et al*. (2006) and the Fourth Assessment Report of the Intergovernmental Panel on Climate Change in Yohe *et al*. (2007) and elsewhere in and (IPCC 2007a, 2007b) make it clear that limiting atmospheric concentrations of greenhouse gases to any specific level cannot guarantee that increases in global mean temperature will be held below any target identified as the boundary of "dangerous" climate change regardless of how it is identified.

7.2 A NEW SOURCE OF CONCERN: WEITZMAN'S DISMAL THEOREM

This debate between the cost-benefit approach and the precautionary approach has recently been informed by a "Dismal Theorem" offered by Weitzman (2007). It shows that profound uncertainty about fundamental parameters like climate sensitivity cannot be overcome for any positive rate of risk aversion and any positive rate of pure time preference for any distribution of events (outcomes) whose moment generating function is infinite and includes the potential for catastrophic climate impacts (here defined as a prolonged period of falling welfare per capita). To be more specific, trouble arises for power-law or lognormal distributions or any distribution with "thick tails" where the probability falls only with a power of the size of the event. In these cases, the impact or consequence of an event can grow exponentially while the probability falls with a power law so that the expected impact becomes unbounded. In practice, the theorem draws its significance from our inability to observe the events in the tails with enough frequency to learn anything useful about relative likelihoods of associated catastrophic consequences. It follows that uncertainty will dominate any calculation of expected climate damage because Bayesian learning about the critical variables (even with very strong time discounting) is never strong enough to keep expected marginal damages finite.

Weitzman's "Dismal Theorem" clearly casts doubt on results derived from a cost-benefit approach to climate policy, at least for studies in which the equity implications of declining marginal utility are recognized. Indeed, Weitzman has suggested that a warning label be attached to integrated assessment models that rely on the cost-benefit approach – something like "Warning: To be applied only to non-extreme climate change possibilities". The Dismal Theorem marginalizes the debate over the social cost of carbon and the associated discussions about what makes estimates high or low because it means that all of the existing

[3] Because tolerable boundaries are typically defined in terms of temperature limits and because temperature change depends, to a first approximation, on cumulative emissions over long periods of time, the appropriate economic response can be visualized by solving for an initial shadow price for carbon (and other warming gases) with the expectation that it would increase over time at an endogenously determined rate of interest.

estimates are infinitely too small. It similarly renders obsolete the current obsession of some of the scientific community for reaching model-based consensus on central tendencies about climate change.[4] The action is, quite simply, in the dismal tails.

On the positive side, the result indicates that the value of some types of information is far greater (and perhaps infinitely greater) than the value of other information. It can therefore offer some guidance on where to devote scarce research resources in climate and policy science. Moreover, it seems to offer sound theoretic footing for a generalized precautionary approach designed explicitly to examine and clarify the definition of tolerable climate change. More careful examination of these implications suggests that another warning label needs to be written, but more on that later.

Before proceeding to make that point, it is important to focus on one important condition of the Dismal Theorem – that decision-makers view the world with some aversion to risk (and thus some aversion to inequality).[5] We could therefore find our way around the Dismal Theorem by simply asserting that policymakers should always proceed as if they were completely risk-neutral. Doing so would, however, mean rewriting much of current economic policy; and doing so only in the climate arena would mean that the United Nations Framework Convention on Climate Change would have to be completely overhauled.

Since neither of these responses will be accepted by the policy community, there is no easy way to dismiss the implications of the Dismal Theorem for climate policy and climate science. To explore these implications a little more fully, it is appropriate to contemplate its applicability in a few different cases. Tol (2003), for example, worked within a cost-benefit framework that recognized multiple regions with and without equity weighting. Even without recognizing the consequences of thick tails in the distribution of climate sensitivity, his Monte Carlo simulations noted the small but non-zero probability that marginal utility could grow infinitely large in one or more regions where even "routine" climate change, particularly when it materializes in the form of declining precipitation, can drive economic activity to subsistence levels. As long as these regions were given non-zero weight in the expected utility calculation, their plight would dominate the policy calculus because expected marginal damages would approach infinity. This was, perhaps, a precursor of the Dismal Theorem.

7.3 IMPLICATIONS OF THE "DISMAL THEOREM"

Yohe (2003) suggested that the problem highlighted in Tol (2003) could be overcome by implementing a second policy instrument designed to maintain economic activity above subsistence levels everywhere – a foreign aid program designed simply to prevent economic collapse anywhere in real time. Tol and Yohe (2007) examined this suggestion within the original modeling framework and found that, with sufficient aid, the issue of infinite marginal damage could be avoided. While this work did not envision events characterized in the fat tails of climate sensitivity, it nonetheless suggests that timely social or

[4] Evidence of this obsession is seen in IPCC (2007a) where the potential contributions of Greenland Ice Sheet melting and collapse of the West Antarctic Ice Sheet from sea level rise estimates were deleted (even though they had been included in IPCC (2001) because there was no model based scientific consensus that could explain what is going on (IPCC, 2007a). In the logic of the Dismal Theorem, this makes the ice sheets more policy relevant, not less.

[5] This assumption is captured simply by allowing the marginal utility of consumption to rise indefinitely as consumption falls to a subsistence level (and to fall as consumption rises beyond the range currently experienced by developed economies).

economic interventions that effectively "lop off the thick tails" of regional climate impacts could undercut the power of the Dismal Theorem. If, however, the impacts of the profound uncertainty were felt globally so that no country or region would have the wherewithal to underwrite the subsistence of another, then the Dismal Theorem could still persist. It is here, therefore, that a generalized precautionary principle – the logical implication of the Dismal Theorem – comes into play.

Can the Dismal Theorem inform the boundaries of precaution? To answer this question, it is important to recognize that these boundaries can be defined in many different ways. Put another way, policymakers are not confined to Bayesian learning about the climate sensitivity and other critical parameters in climate models, and this is a good thing. Roe and Baker (2007) show, for example, that "the probability of large temperature increases" is "relatively insensitive to decreases in uncertainties associated with the underlying climate processes". Allen and Frame (2007) responded by arguing that it was pointless for policy makers to count on narrowing this fundamental uncertainty. Rather than tilting at this (and other similar) windmill(s) like Don Quixote, perhaps the policy community should ask the research community to develop greater understandings of the fundamental processes in other areas – processes that produce catastrophic *impacts* from whatever climate change happens to materialize, for example. Even if they cannot rely on the scientific community to reduce the range of "not implausible" scenarios in the temperature domain, they could ask it to (1) explore the triggers of more regional catastrophe, (2) identify the parameters of fundamental change that define those triggers, (3) contribute to the design of monitoring mechanisms that can track the pace of change relative to these triggers, and (4) conduct small- and large-scale experiments in models, laboratories and perhaps the real world to learn more about the relevant processes. Assuming that the rate of change of these manifestations of climate change could be calibrated to something like the pace of change in global mean temperature, it might then be possible to calibrate some of the fuzzy and politically determined boundaries of "dangerous anthropogenic interference".

Three possibilities emerge for this effort. In the first, regional catastrophic impacts are reversible, but doing so could involve draconian global intervention into the economic sectors from which greenhouse gas emissions were being released. Given the great inertia of the climate and political systems, however, affected societies would probably have to cope with catastrophic impacts for a certain and potentially long period of time. In these cases, the precautionary principle would tell us to restrict emissions along a least cost path for a concentration target as a hedge against both the cost of draconian interventions required to retreat back across the lowest thresholds and the transient costs of enduring "temporary" catastrophes. Nothing would be certain in the calculation of how vigorously to restrict emissions, of course, so the expense involved in their reduction would be have to be seen as an investment in reducing risk – specifically reducing the probability factor in the "probability times consequence" definition of pecuniary risk. Political decisions about exactly how much risk might be considered tolerable would have to be taken, and they would have to evolve as more information about the regional processes became available.

In the second case, one or more of the catastrophic processes is irreversible. Here, the precautionary principle tells us to hedge more strongly against "falling off a cliff". The hedging strategy would presumably impose more stringent emissions reductions much earlier than contemplated in the first instance, and calls for a geo-engineering solution could be expected – a strategy with its own risks, to be sure. In the third case, one or more of the catastrophes is irreversible and unavoidable. In this extreme possibility, preparing for the worst in the affected regions would be the only option, and global mitigation policy might still operate as if one of the other two cases were in force ubiquitously.

To put these three storylines into a "not-implausible" context, consider the collapse of the Atlantic thermohaline circulation (the THC) as an example of a potentially catastrophic event across many parts of the globe. The higher the climate sensitivity, the more likely it becomes and the sooner it might occur. The implications of such a collapse are unknown, particularly in the socio-economic context, but the planet has experienced another climate equilibrium in which it does not exist. Three different explanations of the process by which it might collapse (Keller *et al*., forthcoming) have been advanced, but each would point to its own critical parameter for monitoring. Because we do not know the precise process, we cannot identify the triggering threshold and so we cannot calibrate global policy in terms of an increase in global mean temperature. Schlesinger *et al*. (2005) and Yin *et al*. (2006) have told us, however, that the THC can collapse in a matter of decades once the trigger is pulled and that reversal, if possible, would take as long as a century to achieve.

Clearly, fundamental research into process understanding of circulation dynamics makes more sense in this example than work designed to make marginal changes in the distribution of climate sensitivity. Anticipating progress there, other research could investigate the sensitivity of least cost approaches to hedging strategies to alternative socio-economic futures and the evolution of new scientific knowledge. To be clear, the policy community would find value in this work *only if* the scientific community could clarify (1) the triggering mechanisms, (2) estimate the lag time between the triggers and climatological commitments to crossing the associated thresholds, (3) devise mechanisms for monitoring circulation intensity and other factors with enough precision to inform the likelihood of commitment, and (4) allow statisticians to calculate probabilities of type 1 and type 2 errors along a range of transient futures based on those monitoring exercises. None of these tasks involves Bayesian learning about climate sensitivity. That is reassuring, but none of them is simple either. Faced with an impossibility theorem and persistent uncertainty about climate sensitivity, however, tackling these difficult problems is the lesser of two evils.

7.4 SOME CONCLUDING REMARKS

We have argued that integrated assessment models that rely on a cost-benefit approach to conduce their policy analyses cannot always accommodate profound uncertainties, particularly in the context of persistent thick tails in the distributions of critical parameters like climate sensitivity. It should now be clear why the scientific community must move beyond trying to nail down consensus about the central baseline tendencies of climate change and embrace (though not exclusively) an organized effort designed to examine the "dark tails" of our possible futures across the range of possible impacts and associated key vulnerabilities. Only then can we begin to define the boundaries of tolerable change to support rigorous analyses of decision-making criteria that account, explicitly, for the enormous uncertainties that characterize our understanding of the climate system.

What does all of this mean for the social cost of carbon? Cast in the context of an informed and rigorously defined precautionary approach to policy design, the social cost of carbon can be viewed as the marginal cost of mitigation at any point in time – i.e., the shadow price of the precautionary constraints that reduce the likelihood of catastrophic impacts to tolerable levels. In other words, the calculation of the social cost of carbon would be tied directly to the scarcity rent that minimizes the expected cost of politically palatable hedging. This is not necessarily an easy calculation, but there is some good news. Climate sensitivity would not be an issue because the social cost of carbon would be tied to the marginal cost of meeting a concentration target (though the distribution of climate sensitivity would be involved in the discussions that try to translate temperature targets into concentration limits). The discount rate would not be an issue either, because the rate applied to other public investments and not the one that ponders the ethical complications of intergenerational equity would now apply. Indeed, this calculation would exclude some of the sources of uncertainty that explain the enormous range of social cost of carbon estimates noted above. However, issues like valuation and equity weighting do not go away, as they are essential ingredients to the definition of what constitutes a catastrophic impact.

We hope to have shed some preliminary light on the "So what?" implications of the Dismal Theorem on the design of climate policy and climate research. We now turn to the warning label that we promised. The Dismal Theorem is derived from taking limits, so it is tempting to take its conclusion to its logical extremes. One might, for example, read the Dismal Theorem as saying that the value of some improved information about what might be going on in the thick tail of the climate sensitivity distribution is infinite. If that is so, then we need to do as much as we can to sharpen the climate signal by, for example, burning as much coal as quickly as we can. One might also apply the generalized precautionary principle to all social issues for which there are unfortunate consequences in the fat tails of the distributions of critical variables because expected marginal damages are infinite for all of them. But then, how should we set priorities for distributing the planet's finite resources in the social interest? The economic tradeoffs would simply be undefined. Because neither of these implications is particularly attractive, we offer a concluding warning label on the Dismal Theorem: "Warning: Not to be taken to its logical extreme in application to real world problems."

Acknowledgements

The authors gratefully acknowledge the contributions of Martin Weitzman in the preparation of this chapter, as well as the comments offered by participants in the Climate Change Impacts workshop in Snowmass, Colorado in the summer of 2007. All remaining errors reside with the authors.

7.5 REFERENCES

Alcamo, J. and E. Kreileman (1996). Emissions Scenarios and Global Climate Protection. *Global Environmental Change* **6**, 305–344.

Allen, M. and D. Frame (2007). Abandon the Quest. *Science* **328**, 582–583.

Andronova, N.G. and M.E. Schlesinger (2001). Objective Estimation of the Probability Density Function for Climate Sensitivity. *Journal of Geophysical Research* **106**, 605–622.

Cline, W.R. (1992). *The Economics of Global Warming*. Institute for International Economics: Washington, D.C.

Cline, W.R. (1997). Modelling Economically Efficient Abatement of Greenhouse Gases, in Y. Kaya and K. Yokobori (Eds): *Environment, Energy, and Economy*. United Nations University Press: Tokyo, 99–122.

Cline, W. R. (2004). *Meeting the Challenge of Global Warming*. National Environmental Assessment Institute: Copenhagen.

Downing, T. and P. Watkiss (2003). The Marginal Social Costs of Carbon in Policy Making: Applications, Uncertainty and a Possible Risk Based Approach. Paper presented at the DEFRA International Seminar on the Social Costs of Carbon, July 2003.

Hope, C. (2006). The Marginal Impact of CO2 from PAGE2002: An Integrated Assessment Model Incorporating the IPCC's Five Reasons for Concern. *Integrated Assessment* **6**, 1–16.

Intergovernmental Panel on Climate Change (IPCC) (2007a). *Climate Change 2007: The Science. Contribution of Working Group I to the Fourth Assessment Report*, Cambridge University Press: Cambridge, UK.

Intergovernmental Panel on Climate Change (IPCC) (2007b). *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report*, Cambridge University Press: Cambridge, UK.

Keller, K., G.W. Yohe and M.E. Schlesinger (forthcoming). Managing the Risk of Climate Thresholds: Uncertainties and Information Needs. *Climatic Change*.

Maddison, D.J. (1995). A Cost-Benefit Analysis of Slowing Climate Change. *Energy Policy* **23**, (4/5), 337–346.

Manne, A. and R. Richels (1997). On Stabilizing CO_2 Concentrations – Cost-effective Emission Reduction Strategies. *Environmental Modeling and Assessment* **2**, 251–266.

Nordhaus, W. (1991). To Slow or Not to Slow: The Economics of the Greenhouse Effect. *Economic Journal* **101**, 920–937.

Nordhaus, W.D. (1993). Rolling the "DICE": An Optimal Transition Path for Controlling Greenhouse Gases. *Resource and Energy Economics* **15**, 27–50.

Nordhaus, W.D. (1994). *Managing the Global Commons: The Economics of Climate Change*. The MIT Press: Cambridge, MA.

Nordhaus, W.D. (2007). Critical Assumptions in the Stern Review on Climate Change. *Science* **317**, 201–202.

Nordhaus, W.D. and Z. Yang (1996). RICE: A Regional Dynamic General Equilibrium Model of Optimal Climate-Change Policy. *American Economic Review* **86**: 4, 741–765.

Nordhaus, W.D. and J.G. Boyer (2000). *Warming the World: Economic Models of Global Warming*. The MIT Press: Cambridge, MA.

Roe, G. and M. Baker (2007). Why is Climate Sensitivity so Unpredictable? *Science* **328**, 629–632.

Roughgarden, T. and S.H. Schneider (1999). Climate Change Policy: Quantifying Uncertainties for Damages and Optimal Carbon Taxes. *Energy Policy* **27**, 415–429.

Schlesinger, M.E., J. Yin, G. Yohe, N.G. Andronova, S. Malyshev and B. Li (2005). Assessing the Risk of Collapse of the Atlantic Thermohaline Circulation, in J. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, and G. Yohe (Eds) *Avoiding Dangerous Climate Change: A Scientific Symposium on Stabilisation of Greenhouse Gases.* Cambridge University Press: U.K. Meteorological Office, Exeter, U.K., 37–47.

Stern, N.H., S. Peters, V. Bakhshi, A. Bowen, C. Cameron, S. Catovsky, D. Crane, S. Cruickshank, S. Dietz, N. Edmonson, S.-L. Garbett, L. Hamid, G. Hoffman, D. Ingram, B. Jones, N. Patmore, H. Radcliffe, R. Sathiyarajah, M. Stock, C. Taylor, T. Vernon, H. Wanjie and D. Zenghelis (2006). *The Economics of Climate Change: The Stern Review.* Cambridge University Press, Cambridge, UK.

Swart, R., M. Berk, M. Janssen, E. Kreileman, and R. Leemans (1998). The Safe Landing Approach: Risks and Tradeoffs in Climate Change, in J. Alcamo, R. Leemans, and E. Kreileman (Eds): *Global Change Scenarios of the 21st Century – Results from the IMAGE 2.1 Model*. Pergamon/Elsevier Science: Oxford: 193–218.

Tol, R.S.J. (2002). Welfare Specifications and Optimal Control of Climate Change: An Application of FUND. *Energy Economics* **24**, 367–376.

Tol, R.S.J. (2003). Is the Uncertainty about Climate Change Too Large for Expected Cost-benefit Analysis?. *Climatic Change* **56**, 265–289.

Tol, R.S.J. (2005). The Marginal Damage Costs of Carbon Dioxide Emissions: An Assessment of the Uncertainties. *Energy Policy* **33**, 2064–2074.

Tol, R.S.J. (2007). The Social Cost of Carbon: Trends, Outliers and Catastrophes. Economics Discussion Papers, http://www.economics-ejournal.org/economics/discussionpapers/2007–44.

Tol, R.S.J. and G.W. Yohe (2007). Infinite Uncertainty, Forgotten Feedbacks, and Cost-Benefit Analysis of Climate Change. *Climatic Change* **83**, 429–442.

Toth, F.L., H.-M. Fussel, M. Leimbach, G. Petschel-Held and H.-J. Schellnhuber (1997). The Tolerable Windows Approach to Integrated Assessment, O. K. Cameron, K. Fukuwatari and T. Morita (Eds), *Climate Change and Integrated Assessment Models (IAMs) – Bridging the Gaps*. *Center for Global Environmental Research*, National Institute for Environmental Studies: Tsukuba, Japan, 401–403.

Uzawa, H. (2003). *Economic Theory and Global Warming*. Cambridge University Press: Cambridge, UK.

Weitzman, M.L. (2007). The Role of Uncertainty in the Economics of Catastrophic Climate Change. http://www.economics.harvard.edu/faculty/Weitzman/papers/Catastrophe.pdf.

Wigley, T., R. Richels and J.A. Edmonds (1996). Economic and Environmental Choices in the Stabilization of Atmospheric CO_2 Concentrations. *Nature* **379**, 240–243.

Yin, J., M.E. Schlesinger, N.G. Andronova, S. Malyshev and B. Li (2006). Is a Shutdown of the Thermohaline Circulation Irreversible? *Journal of Geophysical Research*.

Yohe, G.W. (2003). More Trouble for Cost-benefit Analysis. *Climatic Change* **56**, 235–244.

Yohe, G. and Toth, F. (2000). Adaptation and the Guardrail Approach to Tolerable Climate Change. *Climatic Change* **45**, 103–128.

Yohe, G.W., R.D. Lasco, Q.K. Ahmad, N. Arnell, S.J. Cohen, C. Hope, A.C. Janetos and R.T. Perez (2007). Perspectives on Climate Change and Sustainability. M.L. Parry, O.F. Canziani, J.P. Palutikof, C.E. Hanson and P.J. van der Linden (Eds): *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press: Cambridge.

Yohe, G. and D. Tirpak (2008). A Research Agenda to Improve Economic Estimates of the Benefits of Climate Change Policies. *Integrated Assessment Journal* **8**, 1–17.