THE TOLERABLE WINDOWS APPROACH: LESSONS AND LIMITATIONS

An Editorial Comment

1. Introduction

The Tolerable Windows Approach (TWA) described in Petschel-Held et al. (1999), is one representative of a new approach to analyzing global change mitigation policy. The TWA attempts to define the boundaries of tolerable change that might serve as guardrails against catastrophic impacts. It then tries to work ‘backwards’ to see how we might constrain the emission of greenhouse gases to guarantee that those boundaries are never crossed and the associated guardrails are never tested. As such, the TWA joins at least three other analytical tacts that have been exercised over the past decade or so to examine mitigation policy. Each approach has its own strengths and its own weaknesses; but the strengths of some have tended to complement the weaknesses of the others so that their combined contribution to our understanding of global change has been greater than the sum of their individual contributions. Armed with insight into the value of multiple approaches, both the research community and its constituent collection of policy-types have welcomed TWA into the analytical fold. Indeed, the Third Assessment Report (TAR) of the Intergovernmental Panel on Climate Change (IPCC) will include the TWA in several chapters of the Report of Working Group II on ‘Impacts, Adaptation and Vulnerability’ and in its Cross-Cutting Guidance Paper on ‘Frameworks for Decision Making’.

This inclusion is a good thing. The TWA will surely complement the other approaches well. It holds the potential to broaden our understanding of global change and global change policy by focusing attention on issues that might otherwise have been missed. That said, it must be emphasized that the TWA does not have exclusive claim to the truth. Care must be taken to make certain that the approach, taken to its logical extreme, does not lead us to ignore the lessons drawn from other perspectives. This short essay will try to make this point by offering a stylized taxonomy of the collection of analytical approaches to the mitigation policy issue. It will mention several strengths and weaknesses of each. It will recall some lessons learned. It will catalogue gaps and next steps in the research agenda. And it will offer a personal view of where and how the TWA might offer the most value-added.

Climatic Change **41:** 283–295, 1999.

2. A Review of Current Assessment Models

It is difficult to author a review of all of the assessment models that have been created over the past few years to analyze the issues surrounding global change. It would, however, be even more difficult to suggest where the TWA fits into this array without providing such a review. I offer, therefore, an abbreviated overview of recent work; but the reader is warned that this discussion has not been designed to be comprehensive. I try, instead, to place the TWA in what can best be described as a rough sketch of the analytical landscape. The interested reader is referred to Rotmans et al. (1998), to Weyant (1997), to Tol and Fankhauser (1997) or to papers published by the researchers themselves for more thorough coverage of the particulars that I have omitted. The reader is also referred to Schneider (1997) for a discussion of the value judgements that each approach brings to the table.

I begin by highlighting Table I (drawn from Weyant (1997)); it offers a list of many of the major assessment models that are currently being exercised. The first column records the acronym by which each model is known while the second lists its principal developer(s). The first column also provides at least one reference for each model – not necessarily the most recent, but one in which the structure of the model is described in some detail. The last three columns give some insight into those structures. The fourth column highlights differences in the spatial resolution of the twenty models listed there, while the final column indicates how the potential damage associated with climate change is included.

The third column in Table I offers a rough typology, drawn from Weyant (1997), for sorting the listed models. The categories are meant to be suggestive only of the original intent of the modelers, and so they are not really mutually exclusive. Five models are designated DMUM indicating that they were designed most explicitly to accommodate uncertainty. The other fifteen are deemed to be deterministic in nature, indicating that they were designed either to uncover an optimal mitigation policy (DPOM) or to explore the properties of a variety of specific policy options (DPEM). Several of the optimization models have, of course, been employed to conduct second-best analyses of the sort targeted by the creators of evaluation models; and thirteen of the models listed have been employed to conduct uncertainty analysis of one sort or another.

2.1. OPTIMIZATION MODELS

Regardless of their specific application from one context to another, the optimization models listed in Table I were all designed explicitly to characterize mitigation policies that solve a full-blown dynamic problem of maximizing the discounted value of social welfare. They all essentially adopt a cost-benefit approach in which all potential damages are monetized and in which the distributional effects of climate change and climate change policies are essentially ignored. They are all deterministic, as well, so they assume implicitly that current decision-makers have

TABLE I

Integrated assessment models[a]

Model acronym (reference)	Principal developers	Typology[b]	Spacial scale[c]	Damage representation[c]
AS/ExM (Lempert et al., 1994; Lempert et al., 1996)	Rob Lempert/Steve Popper (Rand) Michael Schlesinger (U of Illinois)	DMUM	Global	1 reduced form equation
AIM (Morita et al., 1993)	T. Morita, M. Kainumn (NIES, Japan) Yuzuri Matsuoka (Kyoto U.)	DPEM	$\frac{1}{2}^{\circ} \times 1\frac{1}{2}^{\circ}$ grid	Sectoral models
CETA (Peck and Teisberg, 1992)	Stephen Peck (EPRI) Thomas Teisberg (Teisberg Assoc.)	DPOM*	Global	1 reduced form equation
Connecticut (Yohe and Wallace, 1996)	Gary Yohe (Wesleyan University)	DPOM/ DPEM*	Global	1 reduced form equation
CSERGE (Maddison, 1994)	David Maddison (U. College of London)	DPOM	Global	2 reduced form equations
DIAM (Grubb et al., 1995; Grubb, 1997)	Michael Grubb, M. H. Dong, T. Chapius (Royal Institute of International Affairs)	DPOM*	Global	1 reduced form equation
DICE (Nordhaus, 1994)	William Nordhaus (Yale Univ.)	DPOM*	Global	1 reduced form equation
FUND (Tol, 1995)	Richard Tol (Vrije Universiteit Amsterdam)	DPOM*	Nine regions	Reduced form by sector
HCRA (Hammitt et al., 1995)	Jim Hammitt (Harvard) Atul Jain/Don Wuebbles (U. of Illinois)	DPOM*	One or two Regions	Reduced form with Catastrophes
ICAM-2 (Dowlatabadi and Morgan, 1993)	Hadi Dowlatabadi (Carnegie Mellon) Granger Morgan (Carnegie Mellon)	DMUM	Seven regions	Sectoral models
IMAGE 2.0 (Alcamo, 1994)	Joe Alcamo, M. Janssen, M. Krol (RIVM, Netherlands)	DMEM	$\frac{1}{2}^{\circ}$ grid	Several models
MARIA (Mori, 1995)	Shunsuke Mori (Sci. Univ. of Tokyo)	DPOM	Four regions	1 reduced form equation
MERGE 3.0 (Manne et al., 1995)	Alan Manne (Stanford) Robert Mendelsohn (Yale) Richard Richels (EPRI)	DPOM*	Nine regions	2 reduced form equations
MiniCAM (Edmonds et al., 1994)	Jae Edmonds (Pacific Northwest Lab) Richard Richels (Electric Power Res. Inst.) Tom Wigley (UCAR)	DPOM	Eleven regions	Sectoral models
MIT (Yang et al., 1996)	Henry Jacoby/Ron Prinn (MIT) Zili Yang (MIT)	DPEM	Eleven regions	Sectoral models
PAGE (CEC, 1992; Plambeck and Hope, 1996)	Chris Hope (Cambridge Univ.) John Anderson/Paul Wenman (Env. Res.)	DMUM	Seven regions	Reduced form by sector
PEF (Cohan et al., 1994)	Joel Scheraga/Susan Herrod (EPA) Rob Stafford/Nathan Chan (DFI)	DMUM	Two regions	Sectoral models
RICE (Nordhaus and Yang, 1996)	William Nordhaus (Yale Univ.) Zili Yang (MIT)	DPOM	Six regions	1 reduced form equation
SLICE (Kolstad, 1993; Kolstad, 1994)	Charles Kolstad (Univ. California Santa Barbara)	DMUM	Global	1 reduced form equation
TARGETS (Rotmans et al., 1994)	J. Rotmans (RIVM) J. Janssen (RIVM) H. J. M. de Vries (RIVM)	DPEM*	Varied and detailed	Sectoral models

[a] Adapted from Table 1 and Figure 2 of Weyant (1997).
[b] Notationally, DPOM = Deterministic Policy Optimization model; DPEM = Deliministic Policy Evaluation Model; and DMUM = Decision Making Under Uncertainty Model.
[c] Adapted from Tol and Frankhauser (1997).

a clear understanding of how the future will unfold on both the driving and impact sides of the global climate system.

Optimization exercises in this context generally show what might be appropriate under the best of circumstances as long as potential impacts are gradual, predictable, and reversible. They produce benchmarks of intervention under heroic assumptions about what we know presently – benchmarks of efficiency against which other results can be measured. They serve, therefore, much the same function for climate policy as analyses of perfect competition do in the economics literature. They show, in particular, that only modest emissions control would be warranted over near and medium term policy horizons in the best of circumstances. Most optimal carbon 'taxes' start low ($5 to $15 per ton of carbon in the near-term) and climb slowly over the next century (reaching $50 to $100 per ton by 2100). Why? The models generally relate the cost of emissions reduction to the current emissions rate, and so the costs of emissions reductions begin with the onset of the policy. The impacts of climate change are, meanwhile, usually related to atmospheric concentrations that are modeled to depend upon cumulative emissions given a static biosphere. Current emissions therefore influence the modeled climate system with a lag so that the benefit of the immediate cost of near-term emissions reduction is delayed. These benefits are then discounted because the optimization approach implicitly compares the return to investment in near-term emissions reduction to the returns of other economic investments whose productivity could underwrite more vigorous emissions reductions in the future. Finally, most of the reduced form damage functions included in assessment models are quadratic or cubic; as a result, early damages are relatively small even when they are not discounted.*

2.2. POLICY EVALUATION MODELS

Several of the models listed in Table I (models designated DPEM) were designed explicitly to evaluate the relative efficacy of various policy options, but most of the models listed there have been used at one time or another to support some sort of 'non-optimization' evaluation analysis. The results have sometimes been startling. Wigley et al. (1996), for example, identified (WRE) emissions trajectories that would reduce the discounted cost of holding the concentration of greenhouse gases below specific thresholds by as much as 80% in comparison with the S350-S750 trajectories proposed initially by Working Group I of the IPCC. Others have confirmed their results directly from their models. They all noted the enormous cost-reducing potential of adopting efficient spatial and temporal distributions for

* It cannot be asserted that all optimization results support modest near-term mitigation. Grubb et al. (1995) is representative of a group of dissenting researchers who argue that there is significant inertia in economic infrastructure. Their results therefore show that the requisite technological advances would materialize in the future only if they were induced by relatively restrictive near-term policy.

policies designed to achieve specific environmental quality targets that were *not necessarily chosen as solutions of any dynamic optimization*. Spatial efficiency asserts that the least cost source of emissions reduction should always be exploited, *regardless of its location*. Temporal efficiency, meanwhile, asserts that emissions reduction exploit intertemporal investment opportunities to maximize economic growth given the modeled lags in the climate system. Those opportunities work to 'buy time' before more rigorous emissions reductions are required, but only if emission-saving technologies evolve over the intervening years. Some modelers believe that these technologies can be encouraged even in the absence of near-term mitigation, but others hold that they will *not* emerge if they are not induced by vigorous and virtually immediate intervention.

The so-called 'WRE' scenarios attracted an enormous amount of attention, particularly from those who noted that different emissions trajectories can have different damage profiles even if they hit the same concentration targets in the long run. A productive debate has transpired, but recognizing these differences on the benefit side of the equation can be viewed as a move back toward the cost-benefit framework of dynamic optimization. Recent work in policy evaluation has avoided this regression by examining a number of control-cost/policy-design issues. The advent of the Kyoto Protocol has lead the latest iteration of the Energy Modeling Forum to examine the additional cost involved in meeting politically determined emissions targets for 2012 on the way to long-term concentration targets. Other researchers have begun to explore the relative cost implications of choosing one type of limiting target instead of another (e.g., fixing emissions of greenhouse gases, the associated increase in global mean temperature, or their antecedent atmospheric concentrations, to name but three). Still others have pondered the implications of various emissions permit trading schemes and global allocations on control costs and on the distribution of income. The fundamental lesson here is that policy design matters irrespective of what might happen on the benefit side of a cost-benefit consideration; it is a lesson that should not be lost as the TWA is brought to bear on the problem.

2.3. UNCERTAINTY MODELS

Any consideration of climate change concludes almost immediately that uncertainty is ubiquitous. Uncertainty permeates our understanding of the entire climate system, starting with the drivers of greenhouse gas emissions, cascading through the science of atmospheric interactions of those gases with the oceans and the biosphere, and spilling through the impacts of those interactions on natural, social, political, and economic systems. This observation is obvious, of course, and it is reassuring that it has not been lost on the modeling community. Thirteen models identified in Table I (designated DMUM if they consider uncertainty by design or by an asterisk if their deterministic structures have been applied to uncertain environments) have been exercised in ways that incorporate uncertainty into their analyses to some degree.

The results are, as one would expect, varied; but there are, again, a few general insights to be drawn from the collective efforts of multiple modelers. First of all, anticipating that the globe might be moving along a smooth trajectory characterized by unexpectedly rapid climate change has little effect on the qualitative assessment of dynamically efficient long-term policy. Slightly more rigorous near-term mitigation effort might be appropriate in such a case, but it would still be modest (falling well short of stabilizing emissions at current or historic levels especially if more rapid change were driven by more energetic economic activity). Even near-term hedging against the realization that future change might turn out to be more rapid than expected offers little reason for dramatic intervention over the next several decades.

The story can change significantly, though, if models reflect the possibility that future change might not be smooth. Small chances of high consequence events offer unexpectedly small incentive for adjusting near-term hedging strategies in favor of more robust mitigation, but highlighting their potential for catastrophic and irreversible climate change suggests that policies that satisfy long term dynamic efficiency conditions might be the right answers to the wrong questions. A few analysts have begun to investigate the near-term implications of designing mitigation strategies that can be adjusted over time as we learn more about the size and pace of climate change and its impacts. This has led some to look for robust near-term policies that minimize the expected cost of being wrong during the learning process (see, e.g., Lempert et al. (1996)). It has led others to consider hedging in the near-term against the potential that climate change will be catastrophic before we have a chance to judge precisely what that target should be (see, e.g., Yohe (1997)). Both envision periods of stable policy punctuated by relatively shorter periods of adjustment that can be anticipated by all of the relevant decision-makers.

This work on what Schellnhuber and Yohe (1997) call 'evolutionary assessment' might remind a political economist of 'muddling through'. I find it more satisfying to draw an analogy with how central banks administer the money supply of a developed economy as new information about its health becomes available. The rate of growth in the money supply is typically augmented during periods of downturn. Interest rates fall, as a result, so that investment, aggregate demand and finally employment can all grow. The rate of growth is, of course, correspondingly restrained during inflationary periods. In this case, interest rates climb so that aggregate demand falls and the pace of inflation slows. Indeed, one of the primary functions of a banking system is to maintain a structure through which these policy adjustments can be transmitted effectively throughout an economy in a way that is thoroughly understood and generally foreseen.

How can we apply the monetary policy analogy to climate policy? Here there are only questions to report. We must first determine how we will learn that the climate is changing. How should we set near-term policy to maximize the pace of that learning? When might we learn that future climate change will or will not be smooth when there is so much natural variability in the climate system? When will we be more certain that climate change will proceed rapidly, or slowly, or variably into the future? What sort of monitoring mechanisms should be created so that policy can respond to new understanding that will always be clouded by noise? How do we set climate policy targets that are meaningful, feasible, and flexible? Should the targets be specified in emissions, concentrations, temperature change, the rate of temperature change, or what? How do we create mechanisms so that we can minimize the cost of meeting those targets *and* minimize the cost of making adjustments as the future unfolds? How sensitive are these costs to the specification of the targets? How much short-term variability should be allowed as the globe moves to meet interim policy targets along long-term trajectories? What institutions are required to administer a global policy that addresses these and other issues?

3. Casting the TWA in the Larger Context of Assessment Methods

My impression of the TWA is informed by my participation in several workshops hosted by the Potsdam Institute for Climate Impact Research over the past few years. I begin this section with a brief description of that view. It seems to me that the TWA sets out to define a target-window for mitigation policy that can best be viewed as the intersection of windows of tolerable change drawn from a multitude of climate-sensitive systems. Some of the systems can be economic for which the boundaries of tolerable change can be defined in terms of holding monetized damages below some threshold. But other systems could be social, or political, or ecological; and their boundaries of tolerable change can be defined in other terms. Demographers and sociologists might propose boundaries based on limiting climate-induced migration or social stress. Ecologists might propose boundaries based on limiting ecosystem migration or ecosystem stress. Political scientists might propose boundaries based on preserving certain institutional structures or initiatives. The list is virtually endless – and this is a strength of the TWA.

Figure 1 displays this impression graphically in the simplest case. It depicts boundaries of tolerable change from five systems (labeled A through E) without reference to their sources. Each depicts a threshold of temperature change and/or rate of temperature change that somebody has determined to be the limit of tolerability so that anything to the left and/or below the threshold is deemed to be tolerable. The intersection of these areas of tolerable change [defined from top left to bottom right by Constraints E, D and finally C] represents the target window for mitigation policy *under the assumption that the global decision-making process has accepted all of the underlying boundaries.* Recognizing this underlying structure is, in my view, critical in assessing the role that TWA can play in advancing

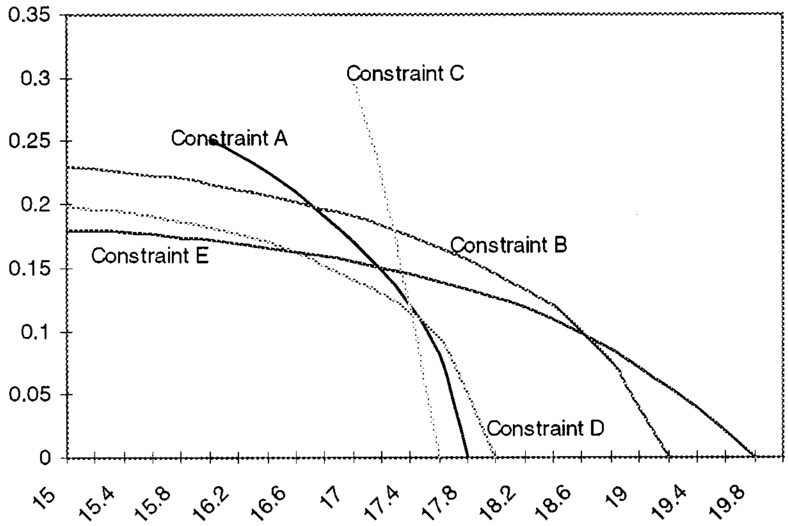

Figure 1. A representation of the underlying structure of a tolerable window derived from five underlying constraints.

our understanding of both clarifying the relevant policy questions *and* defining an efficient research agenda.

3.1. AN EARLY CAVEAT

Before speaking to the many ways in which this view of the TWA illustrates its potential value, it is imperative to underscore why we need to avoid taking this construction to its logical extreme. This is a simple point. If the target window defined in Figure 1 were taken as absolutely binding, then prescribing mitigation policy that would guarantee that we never move beyond its boundary would be equivalent to solving a dynamic optimization problem with zero damage within the window and infinite damage outside. This is not an appropriate interpretation of the TWA. Why? Because we cannot define any of the underlying boundaries with the knife-edge certainty required to sustain such an interpretation. Because to do so would mean that we missed the chance of assessing the opportunity costs of adding increasingly restrictive constraints to the construction of the target window. And because to do so would mean that we had ignored the tradeoff between investing incrementally in tolerable climate change and investing in other things. In short, because to do so would mean that we missed the opportunity to exploit the full power of the TWA in better informing our consideration of mitigation policy.

3.2. SOME OBVIOUS STRENGTHS

Several sources of TWA's power should be obvious from its structure even without going to the extreme. First of all, the TWA does *not* require that all damages be monetized. Many if not most of the underlying constraints could emerge from considerations that are non-economic, and they can carry as much or more weight than objectives drawn from the calculus of economic efficiency. Nor is it necessary to ignore distributional issues. The underlying constraints can easily be identified, and so at least a qualitative accounting of who faces a potential burden from climate change and who does not can easily be sustained. Simply put, some constraints are binding (Constraints C, D, and E, for example, in Figure 1), and so their constituents are at risk; others are not binding (Constraints A and B) contingent on tighter constraints being sustained; and so their constituents are not at risk – at least not yet. Policy directed at staying within the target area would thus explicitly recognize the claims on global policy exercised by the proponents of at least one of the binding constraints; and omitting that constraint would allow policy-makers to compute the opportunity cost of living by that recognition.

3.3. ECONOMICS AND DISCOUNTING

This last observation means that economics and discounting continue to have a role to play in exercising the TWA. It should be clear that the TWA really functions to set the stage for a more sophisticated set of policy evaluation exercises defined across a wider array of systematically chosen objectives; and achieving any objective at least cost is surely an important goal. Minimizing cost can improve the likelihood of implementation, and discounting becomes far less controversial when applied only to costs. Indeed, comparing the discounted cost of achieving various climate objectives with the cost of investing in other welfare enhancing initiatives would be one hallmark of a social and political decision-making process that strives to conserve the globe's finite resources. One could even envision evaluating the incremental cost saving of removing selected underlying constraints from the definition of the target window (e.g., Constraint C in Figure 1) to determine the opportunity cost of including them in the definition of tolerable change.

It turns out that this is not as simple a task as it might appear. Certainly, the opportunity cost of a non-binding constraint like A or B would be zero, but there is more to the story. Figure 2 shows why by superimposing a typical trajectory of temperature and the rate of temperature change that might be observed along an emissions trajectory that has been constrained by the designated window. Notice that it intersects the boundary of the window in two places. One intersection is typical of all paths; the lower right hand corner of any window must come into play over the long term when temperature has reached its maximum so that its rate of change must converge to zero. The other binds in the short term at some positive rate of change and a temperature that is smaller than the maximum allowed. As drawn, this second intersection has been defined by Constraint D, but that need not be the case. Yohe (1997) has shown that the precise location of this early intersection depends critically upon the unregulated path of emissions against which mitigation has been applied. For some trajectories, then, Constraint C might bind

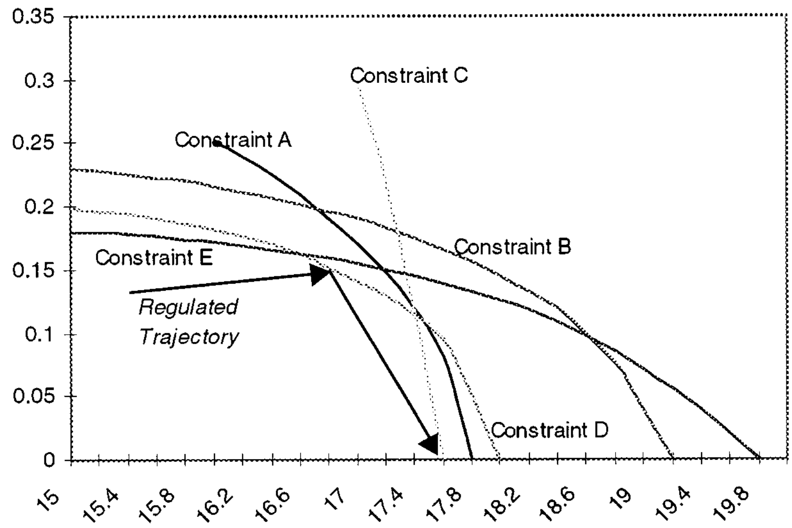

Figure 2. A representation of a regulated trajectory that is constrained by the window first along Constraint D and then at the lower right hand corner where the maximum tolerable temperature is sustained indefinitely.

emissions in the short-term; and for others, Constraint E. The opportunity cost of imposing any underlying constraint on the larger system cannot, therefore, be computed without making some reference to the anticipated baseline trajectory of unregulated emissions.

3.4. ADAPTATION AND UNCERTAINTY

The last section spoke of removing constraints without reference to why policy-makers might want to do so. One obvious reason might be that they turn out to be too expensive when the definition of the window is actually negotiated. Framers of policy in this context need not conduct cost-benefit analyses for each constraint; but they should have some idea of the value (economic, social, political, ecological, ethical, etc.) of maintaining a constraint. Computing the cost of allowing it into the definition of the target window so that it might bind emissions can then offer them the possibility of judging whether or not preserving that value is worth the effort.

The last section also spoke as if the constraints were known with certainty, and this cannot be so. Our understanding of each will surely be fraught with uncertainty – of both the physical impacts *and* the potential of reducing their harm by adapting. It is, in this context, that the structure upon which the TWA is constructed could pay the largest dividends by identifying precisely where improved understanding of climate impacts might pay the largest dividends. We should, in particular, work the hardest to understand the constraints that are most likely to bind: a Constraint like D in Figure 2, but perhaps Constraint E if unregulated emissions track higher than expected so that the rate of change constraint binds earlier. Improved understanding of a previously underestimated adaptive potential could loosen a constraint like D to something like A and thereby diminish the need for restrictive emissions policy in the short-term. Improved understanding of the uncertainties might mean that the potential catastrophic impact upon which a constraint like D was formulated was more or less likely and/or avoidable by other means. Information like this would, of course, strengthen or weaken the resolve of the negotiator to maintain it in the definition of the target window when its significance is compared against its cost.

4. Some Conclusions

To summarize, the Tolerable Windows Approach has many strengths that can be exploited to further our understanding of how to respond to a changing climate. It does not require that all impacts be monetized. It does not ignore distributional issues. It can be used to judge the opportunity cost of imposing increasingly restrictive targets on greenhouse gas emissions; and it can associate these incremental costs with specific sources of concern. It can accommodate thresholds and catastrophes in a way that supports specific concern about specific events without pondering the difficult multiplication of a small probability with a large damage estimate. It can focus research attention designed to facilitate adaptation on areas where adaptation might be most important; and it can focus impacts research on areas where reducing uncertainty can be most valuable. It can help us investigate the sensitivity of these areas of focus to changes in anticipated emissions. It can, in short, support the definition of targets for climate change policy that will move and evolve efficiently over time.

Where does it fall short? It says nothing about how to build and sustain the institutions within which policy-makers will negotiate which constraints are to be allowed and which are not. It says nothing about how to design the mitigation policy that will most efficiently maintain tolerable climate change. It says nothing about how to build and sustain the institutions with which the globe will adjust this policy as our understanding of tolerable change evolves. But it does stand ready to complement other analytical approaches in informing these processes.

Acknowledgements

This work was supported by the National Science Foundation under Cooperative Agreement SBR-9521914. The author thanks Mark Jacobsen and Hadi Dowlatabadi for comments on an earlier draft; remaining errors reside, as usual, with the author.

References

Alcamo, J. (ed.): 1994, *Image 2.0: Integrated Modeling of Global Climate Change*, Kluwer Academic Publishers, Dordrecht.

Cohan, D., Stafford, R. K., Scheraga, J. D., and Herrod, S.: 1994, The Global Climate Policy Evaluation Framework', in the Proceedings of the 1994 A&WMA Global Climate Change Conference, Phoenix, April 5–8, Air & Waste Management Association, Pittsburgh.

Commission of the European Communities [CEC]: 1992, 'PAGE User Manual', Brussels.

Dowlatabadi, H. and Morgan, M. G.: 1993, 'A Model Framework for Integrated Assessment of the Climate Problem', *Energy Policy* **21**, 209–221.

Edmonds, J.A., Pitcher, H., Rosenberg, N., and Wigley, T.: 1994, 'Design for the Global Change Assessment Model', in Kaya, Y., Nakicenovic, N., Nordhaus, W. D., and Toth, R. L. (eds.), *Costs, Impacts and Benefits of* CO_2 *Mitigation*, IIASA, Laxenburg.

Grubb, M., Duong, M. H., and Chapius, T.: 1995, 'The Economics of Changing Course', *Energy Policy* **23**, 417–432.

Grubb, M.: 1997, 'Technologies, Energy Systems and the Timing of CO_2 Emissions Abatement', *Energy Policy* **25**, 159–172.

Hammitt, J. K. and Adams, J. L.: 1995, 'The Value of International Cooperation in Abating Climate Change', American Economics Association/Association of Environmental and Resource Economists Annual Meeting, Washington, D.C., January.

Kann, A.: 1997, 'A Comparison of Approaches for Performing Uncertainty Analysis in Integrated Assessment Models', Energy Modeling Forum, Stanford University, June.

Kolstad, C. D.: 1993, 'Looking vs. Leaping: The Timing of CO_2Control in the Face of Uncertainty and Learning', in Y. Kaya, Nakicenovic, N, Nordhaus, W. D., and Toth, R. L. (eds.), *Costs, Impacts, and Benefits of* CO_2 *Mitigation*, IIASA, Laxenburg.

Kolstad, C. D.: 1994, 'George Bush vs. Al Gore, Irreversibilities in Greenhouse Gas Accumulation and Emission Control Investment', *Energy Policy* **22**, 76–92.

Lempert, R. J., Schlesinger, M. E., and Hammitt J. K.: 1994, 'The Impact of Potential Abrupt Climate changes on Near-Term Policy Choices', *Clim. Change* **26**, 351–376.

Lempert, R. J., Schlesinger, M. E., and Bankes, S. C.: 1996, 'When We Don't Know the Costs or the Benefits: Adaptive Strategies for Abating Climate Change', *Clim. Change* **28**, 235–274.

Maddison, D.: 1994, 'A Cost-Benefit Analysis of Slowing Climate Change', Center for Social and Economic Research on the Global Environment (CSERGE), University College of London, May.

Manne, A. S., Mendelsohn, R, and Richels, R. G.: 1995, 'MERGE: A Model for Evaluating Regional and Global Effects of GHG Reduction Policies', *Energy Policy* **23**, 17–34.

Mori, S.: 1995, 'A Long-Term Evaluation of Nuclear Power by Extended DICE+e Model Simulations-Multi Regional Approach for Resource and Industry Allocation', *Progress Nuclear Energy* **29**, 135–142.

Morita, T., Matsuoka, Y., Kainuma, M., Harasawa, H., and Kai, K.: 1993, 'AIM–Asian-Pacific Integrated Model for Evaluating Policy Options to Reduce GHG Emissions and Global Warming Impacts', National Institute of Environmental Studies, Tsukuba, Ibaraki, 305 Japan, September.

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

Nordhaus, W. D. and Yang, Z.: 1996, 'A Regional Dynamic General Equilibrium Model of Alternative Climate-Change Strategies', *Am. Econ. Rev.* **86**, 741–765.

Peck, S. C. and Teisberg, T. J.: 1992, 'CETA: A Model for Carbon Emissions Trajectory Assessment', *Energy J.* **13**, 55–77.

Petschel-Held, G., Schellnhuber, H.-J., Bruckner, T., Tóth, F., and Hasselmann, K.: 1999, 'The Tolerable Windows Approach: Theoretical and Methodological Foundations', *Clim. Change* **41**, 303–331.

Plambeck, E. L. and Hope, C.: 1996, 'An Updated Valuation of the Impacts of Global Warming', *Energy Policy* **24**, 116–131.
Rotmans, J., van Asselt, M. B. A., deBruin, A. J., den Ellen, M. G. J., deGreff, J., Hilderink, H., Hoekstra, A. Y., Janssen, M. A., Koster, H. W., Martens, W. J. M., Niessen, L. W., and deVries, H. J. M.: 1994, 'Global Change and Sustainable Development: A Modeling Perspective for the Next Decade', Rijksinstituut voor Volksgezondheid en Milieuhygiene, National Institute of Public Health and Environmental Protection, Bilthoven, The Netherlands, June.
Rotmans, J., Dowlatabadi, H., Fialr, J. A., and Parson, E. A.: 1998, 'Integrated Assessment of Climate Change: Evaluation of Methods and Strategies', in *Human Choice and Climate Change*, Chapter 10, Battelle Press, Columbus, OH.
Schneider, S.: 1997, 'Integrated Assessment Modeling of Global Climate Change: Transparent Rational Tool for Policy Making or Opaque Screen Hiding Value-Laden Assumptions?', *Environ. Model. Assess.* **2**, 229–249.
Schellnhuber, J-H. and Yohe, G.: 1998, 'Comprehending the Economic and Social Dimensions of Climate Change by Integrated Assessment', in Proceedings of the Conference on the World Climate Research Programme, Geneva, World Meteorological Organization.
Tol, R. S. J.: 1995, 'The Climate Fund: Sensitivity, Uncertainty, and Robustness Analyses', Working Paper W-95/02, Institute for Environmental Studies, Vrije Universiteit, Amsterdam.
Tol, R. S. J. and Fankhauser, S.: 1997, 'On the Representation of Impact in Integrated Assessment Models of Climate Change', Vrije Universiteit, Amsterdam, April.
Weyant, J.: 1997, 'Insights for Integrated Assessment', IPCC Asia-Pacific Workshop on Integrated Assessment Models, Tokyo, March 10–12.
Wigley, T. M. L, Richels, R., and Edmonds, J. A.: 1996, 'Economic and Environmental Choices in the Stabilization of Atmospheric CO_2 Concentrations', *Nature* **379**, 240–243.
Yang, Z., Eckhaus, R. S., Ellerman, A. D., Fitzmaurice, J, and Jacoby, H. D.: 1996, 'The MIT Emissions Prediction and Policy Analysis (EPPA) Model', Report No. 6, MIT Joint Program on the Science and Policy of Global Change, Massachusetts Institute of Technology, Cambridge, MA.
Yohe, G. and Wallace, R.: 1996, 'Near Term Mitigation Policy for Global Change under Uncertainty: Minimizing the Expected Cost of Meeting Unknown Concentration Thresholds', *Environ. Model. Assess.* **1**, 47–57.
Yohe, G. W.: 1997, 'Uncertainty, Short Term Hedging and the Tolerable Window Approach', *Global Environ. Change* **4**, 303–316.

Department of Economics, GARY W. YOHE
Wesleyan University, Middletown, CT, 06459, U.S.A.

and

Center for Integrated Study of the Human Dimensions of Global Change,
Carnegie Mellon University,
Pittsburgh, 15213, U.S.A.