Pergamon

Global Environmental Change, Vol. 7, No. 4, pp. 303–315, 1997

Printed in Great Britain
0959-3780/97 $17.00 + 0.00

PII: S0959-3780(97)00014-9

Uncertainty, short-term hedging and the tolerable window approach

Gary W. Yohe

Hedging strategies that minimize the expected discounted cost of restricting both temperature change and the rate of temperature change within 'tolerable windows' defined by German researchers are explored. Targeting the smaller window set by the German Advisory Board on Climate Change through the year 2020 emerges as the least-cost strategy as long as the likelihood that this smaller target will ultimately be chosen is greater than roughly 11%. At this break-even likelihood, though, the expected cost, even discounting according to the Ramsey optimal growth rule, amounts to something on the order of $21 trillion (1990 $). Since it is the rate of change constraint that binds more immediately, this estimate can be interpreted as the opportunity cost of limiting temperature change to no more than 0.2°C through the year 2020.

Department of Economics, Wesleyan University, Middletown, CT 06459, USA. Center for Integrated Study of the Human Dimensions of Global Change, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Literally dozens of integrated assessment models of global climate change have been created over the past few years by researchers scattered all over the world.[1] Each model has its own strengths and weaknesses because each has been constructed to focus on one or another aspect of how best to respond to the potential damage of climate change, the likely cost of mitigating its driving forces, or both. Some models are very detailed and very disaggregated. They are designed specifically to ponder the sectoral or regional impacts of change, adaptation to those impacts, appropriate mitigation policy, or adaptation to that policy. They are elaborate, complicated, sometimes difficult to understand, always expensive to create and to maintain, and frequently troublesome to exercise even for a specific and sometimes relatively narrow set of questions. Other models are very simple and highly aggregated. In contrast to their large counterparts, they are designed to examine uncertainty, the value of information, and the potential that new information or new approaches to the policy issue might actually support significant change in the appropriate global response. They are generally more easily understood, versatile, suggestive in their results, relatively inexpensive to create, and frequently easy to exercise even over a wide range of possible futures.

This paper hopes to show the value of creating and maintaining the latter by reporting the results of exercising one simple model designed explicitly to accommodate uncertainty in investigating the potential role of hedging in a new context: setting global policy when the value of mitigation is cast in terms of staying within a 'tolerable window,' of global change. The notion is simply to write global policy so that future conditions stay within a 'window' that places 'tolerable' limits on the pace and level of temperature change over the very long term. It must be emphasized at the beginning that care will need to be taken in interpreting these results. Global-scale models never produce quantitative descriptions of policy that can be believed beyond one or two significant figures, and the results reported here are no exception. Even when their application is most appropriate, simple aggregate models produce qualitative insight into first-order effects; and so they produce, at best, hypotheses whose content should be expressed in well defined and researchable ques-

[1]See Rotmans and Dowlatabadi (1996), documentation of EMF-14 activity (1994) or Weyant (1997) for descriptions of many of the existing models.

tions and whose robustness should be tested within the more detailed environments of the more complicated models to those questions.

The results reported here do, in fact, offer a hypothesis that may be worthy of some further examination by researchers who come to the table equipped with more intricate models. More specifically, explicit consideration of hedging against suffering the extreme long-term cost of staying within a tightly framed 'window' that jointly constrains the pace and level of greenhouse-induced temperature change seems to support undertaking unusually severe short-term mitigation. This hypothesis draws support from an economic analysis despite the obvious cost of near-term mitigation and so it runs contrary to results derived in many previous examinations of hedging against bad outcomes.[2] The key difference, it turns out, is the rate of change constraint. Indeed, the high cost of the near-term emissions reduction that finds some justification here can be viewed as the potential opportunity cost of restricting the rate of change of global mean temperature to no more than 0.2°C per decade through the year 2100 and beyond.

Section 1 briefly describes two alternative windows whose temperature boundaries have recently emerged from researchers in Germany. Section 2 highlights how emissions might be restricted over the long term along five alternative and probabilistically weighted trajectories so that future temperature change fits within those boundaries; i.e., the policies drafted with perfect information in 1997 are portrayed. Section 3 reports on the results of the hedging exercise. Alternative emission reduction strategies, each the solution of a deterministic problem in Section 2, are considered in the short run when it is unclear which emissions trajectory will actually emerge and which of two tolerable windows will be chosen as the long-run policy objective. It will be assumed that all is made clear sometime after the year 2010 so that long-term mitigation policy can adjust fully to that revelation by the year 2020. Section 4 finally casts the necessarily quantitative discussion of Section 3 in terms of insight, hypothesis, specific inquiries worthy of subsequent investigation, and lessons for the long-term global change research agenda.

1. Tolerable window targets

The present analysis is not a cost–benefit analysis of mitigation policy. It builds, instead, from the assumption that future mitigation policy *will* be designed to steer the future into one of two alternative 'windows' that define the limits of 'tolerable' global climate change in terms of temperature change and the rate of temperature change. It asks not whether setting such targets is appropriate, but how much it might cost to keep future temperature change trajectories within one or the other target window; and it concludes by pondering the near-term implications of waiting twenty or so years before deciding exactly which of the two alternative targets will actually be chosen.

It should be noted that the tolerable windows are, themselves, derived independently of the cost that might be involved in meeting their implicit limits on greenhouse gas emissions. The tolerable windows approach is, instead, based on "external normative specifications of tolerable sets of climate impacts" (Toth, 1997). Each window is deduced by analyzing the maximum climatic stress that can be assumed to be ecologically and socio-economically bearable; and the problem of determining admissible emissions paths that can meet the limits imposed by those maxima is

[2]The results of the Uncertainty Working Group of EMF-14 generally show that hedging even against very high consequence (but low probability) events adds but a small risk premium to the trajectory of economically efficient carbon taxes. See Yohe (1996) and Manne (1996).

confronted only after those limits are defined. The German Advisory Board on Global Change (WBGU), for example, defined a window based on criteria derived from the earth's geological history and from other considerations [see WBGU (1995)]. Theirs is the smaller window depicted in Figure 1. Denoted henceforth by 'WBGU', this smaller window would (1) limit total temperature change beyond the year 1997 to 1.5°C and (2) restrict the rate of change over time to 0.2°C per decade as long as it does not become more than 0.5°C warmer. Beyond 0.5°C, the maximum tolerable rate of change falls quadratically until it reaches zero at a maximum tolerable warming of 1.5°C. The second window portrayed in Figure 1 will be denoted 'Larger' in what follows. It was adopted from Toth (1997) as an alternative, less severe policy target. The Larger Window is simply wider than its WBGU alternative. The rate of temperature change is still limited to a maximum of 0.2°C per decade, but the quadratic decline in the maximum rate of change does not begin until it becomes 1.5°C warmer. Notice that the limiting rate of change does not reach zero until the temperature climbs 2.5°C higher than current conditions.

Figure 1 also displays combinations of temperature change and rate of change that can be achieved for the median emissions trajectory defined in Yohe and Wallace (1996). Unconstrained, the global mean temperature can be expected to grow along that emissions path by an additional 3.33°C through the end of the next century; and the rate of temperature change peaks at 0.223°C per decade early in its time trajectory. Policy can reduce both, of course; but it takes considerable effort to get within either target window. Indeed, as will be reported below, both temperature and the rate of change can be achieved by lowering emissions until the Larger and WBGU Windows are achieved at (1.74°C, 0.191°C per decade) and (1.04°C, 0.141°C per decade), respectively. Meeting the boundary of the Larger Window requires imposing a carbon tax of \$8.98 per ton in 1997 and allowing that tax to climb (roughly exponentially) to approximately \$235 by the year 2100; atmospheric concentrations of carbon are, in achieving the Larger Window, effectively constrained to less than 540 ppm. Meeting the boundary of the WBGU Window is much more expensive; the carbon tax, in this case, must climb from \$40.15 per ton in 1997 to nearly \$960 by 2100, and concentra-

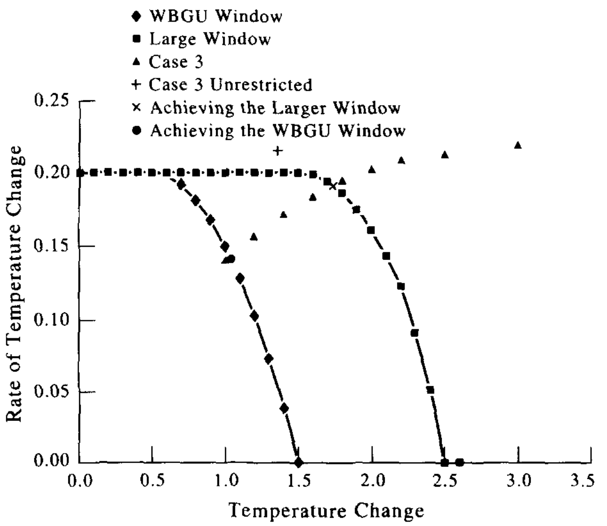

Figure 1. Least cost combinations of change in global mean temperature and the rate of temperature change that can be achieved along the median emissions trajectory are compared with the boundaries of the WBGU and Large tolerable windows.

tions are limited to less than 440 ppm. The former reduces cumulative emissions through 2100 by 34.3% with discounted control costs of \$2.25 trillion;[3] the latter, by 60.9% with discounted control costs totaling \$17.43 trillion.[4]

2. Achieving the tolerable windows — deterministic analysis

The median trajectory noted above was drawn from a distribution of possible emissions futures generated from a simple, aggregate, and probabilistically exercised model of global economic activity, documented in Yohe and Wallace (1996).[5] It is, in fact, most accurately described as the median of a discretized representation of that distribution. As such, it is one of five subjectively weighted emissions alternatives that were chosen to reflect most accurately the full range of possible emissions trajectories supported by the model.[6] Table 1 demonstrates that, taken together, they nearly span the emissions range for the year 2100 of the six IS92 scenarios created by the IPCC (1992), and they straddle more than 90% of the range of modeler's choice runs reported to EW-14 (1994)[7]. Notice that the case 3 trajectory coincides well with the IPCC 'business as usual' IS92a scenario. Case 1 does not limit emissions as severely as IS92c, but it is the result of a fortunate combination of driving variables and behavioral parameterization that does not depend upon any of the restrictive policy assumptions that define IS92c. On the high side, case 5 runs above IS92e; but it is representative of almost 20% of the runs that emerge from the Yohe and Wallace (1996) Monte Carlo analysis of potential futures. Coincidence in reported ranges of

Table 1 Comparisons of the representative scenarios with IPCC IS92 scenarios and EMF-14 alternatives

Case[a]	Representative scenarios: Emissions in 2100[b]	Concentrations in 2100[c]
1	7.8	502
2	15.6	616
3	20.2	679
4	28.7	785
5	43.4	972
Case	**IPCC IS92 scenarios:[d] Emissions in 2100**	**Concentrations in 2100**
c	4.6	n/a
d	9.9	n/a
b	18.6	n/a
a	19.8	n/a
f	25.9	n/a
e	34.9	n/a
Case	**EMF-14 representative scenarios:[e] Emissions in 2100**	**Concentrations in 2100**
modeler's choice (L)	8.5	605
modeler's choice (H)	32.0	1150
standardized (L)	12.0	610
standardized (H)	48.2	1550

[a]Represenatiave scenarios drawn from Yohe and Wallace (1996); they have subjective probability weights of 0.27, 0.13, 0.23, 0.19 and 0.18, respectively.
[b]Measured in billions of metric tons of carbon.
[c]Measured in parts per million by volume.
[d]Six afternative scenarios reported in IPCC (1992). Each has its own underlying description and much of the variation is derived from alternative assumptions about energy and environmental policy. [e]Values drawn from the second round of results reported to the Energy Modeling Forum (EMF-14) (1994) by a variety of modelers; (L) and (H) represent the lowest and highest reported results in emissions and/or concentrations. The modeler's choice runs are derived using the individual modeler's expectations of the trajectories of major driving variable; the standardized runs are derived from a common set of expectations.

[3]All costs quoted in this paper are in 1990 US\$.

[4]Control cost are measured in terms of the dead weight loss created by any carbon tax. The statistics reported here and throughout the paper are those losses discounted at the Ramsey rate (the sum of an assumed 3% rate of pure time preference and the endogenously determined annual rale of growth of per capita consumption) and summed through the year 2100. Implicit here and explicit in the model is a social welfare function that is logarithmic in per capita consumption. See Koopmans (1967).

[5]The model is a marriage of Nordhaus and Yohe (1983) and Nordhaus's DICE (1994). It is an aggregate model with global output driven by four inputs (capital, labor, fossil-based energy and non-fossil energy) whose potential for substitution is quantified by elasticities of substitution between the two types of energy *and* between energy and the other inputs. The supply of fossil fuel is determined in any period by resource availability and the cumulative effects of technological change; the supply of non-fossil fuel is determined by a bias in technological change into or away from this alternative. And the link to the natural climatic system is provided by a simply marginal airborne fraction model coupled with a two-equation temperature module.

continued on page 307

Table 2 Summary statistics for achieving tolerable windows along representative scenarios with perfect information

Case[a]	Emissions reduction[b]	Control cost[c]	Max. temp.[d]	Max. rate[e]
Achieving the WBGU Window (Increase in mean global temperature ≤ 1.5°C):				
1	32%	1.2	1.04	0.143
2	55.30%	5	1.02	0.146
3	60.90%	17.4	1.04	0.141
4	67.70%	24.3	1.07	0.144
5	78.80%	300.4	1.05	0.136
Achieving the Larger Window:				
Case[a]	**Emissions reduction[b]**	**Control cost[c]**	**Max. temp.[d]**	**Max. rate[e]**
1	0%	0	1.56	0.173
2	23.40%	0.4	1.76	0.189
3	34.20%	2.2	1.74	0.191
4	47.10%	5	1.72	0.192
5	61.20%	48.9	1.84	0.182

[a]Identification of the unregulated emissions trajectories ordered from lowest to highest. Table 1 casts each against the IPCC IS92 scenarios as well as the low and high modeler's and standardized runs from EMF-14.
[b]Measured in terms of cumulative emissions through the year 2100 against the unregulated trajectories for each case.
[c]Measured in terms of the present value of total deadweight loss reflected under the derived demand for fossil fuel through the year 2100 and denominated in trillions of 1990 dollars. The discount rate is the Ramsey rate: 3% plus the contemporaneous rate of growth of per capita consumption.
[d]Maximum increase in global mean temperature from the 1990 level.
[e]Maximum increase per decade in global mean temperature from the year 2000 forward

outcomes is not validation, to be sure; but it is reassuring for an exploratory exercise that most of what published work offers for descriptions of the long-term future has been captured by the distribution of outcomes generated by the underlying model.

Table 2 summarizes what would be involved in achieving the WBGU and Larger Windows along the five trajectories, assuming that we know in 1997 which most closely describes how the future will unfold. That is to say, both panels describe policies that would emerge from an unrealistically deterministic analysis of the cost of keeping the future within either window. The five cases are ordered from lowest to highest by the level of unrestricted carbon emissions in the year 2100. Notice, as should be expected, that the reduction in cumulative emissions through the year 2100 would be lower if the Larger Window were targeted. So, too, would control costs. It is, quite simply, easier to keep the trajectories of future temperature change within the Larger Window. Indeed, no control would required if the case 1 trajectory were to materialize because unrestricted emissions would then produce a 'tolerable' temperature profile over time.

It should be noted that the emissions reduction policies whose properties are reported in Table 2 are not 'optimal' in the strict sense of the word. They are not precise solutions derived from a full-blown dynamic optimization approach. They are, however, reasonably accurate reflections of solutions to a conceptual dynamic problem that views the issue of constraining temperature change through the year 2100 along any emissions trajectory as one of holding cumulative emissions through the end of the next century beneath some well defined limit. The policy problem of limiting temperature and its rate of change is thereby converted to one of allocating the fixed quantity of a non-renewable 'resource' (namely carbon emissions) over a fixed (admittedly long) time period, and its solution reduces (roughly) to computing a scarcity rent for 1997 and prescribes a shadow price for carbon that grows exponentially as the future unfolds at the rate of interest. Figure 2A portrays

continued from page 306
[6]The selection process is documented in Yohe (1996). The five cases listed in Table 1 were selected to minimize the sum ofthe squared error of representing the full distribution of emissions in 2100 by five representative scenarios. The selection process also produced subjective relative likelihoods for each: 0.27, 0.13, 0.23, 0.19, and 0.18, respectively.
[7]EMF-14 denotes the 14th model comparison exercise conducted by the Energy Modeling Forum located at Stanford University. Following the lead of EMF-12, its focus has been climate change; and one of its products was to produce distributions of emissions and concentrations scenarios drawn from participating modeling groups. The 'modeler's choice' runs allow each group to define its own set of driving variables; the 'standardized' runs depict results from the same modelers when they adopt the EMF-14 standard profiles for those variables.

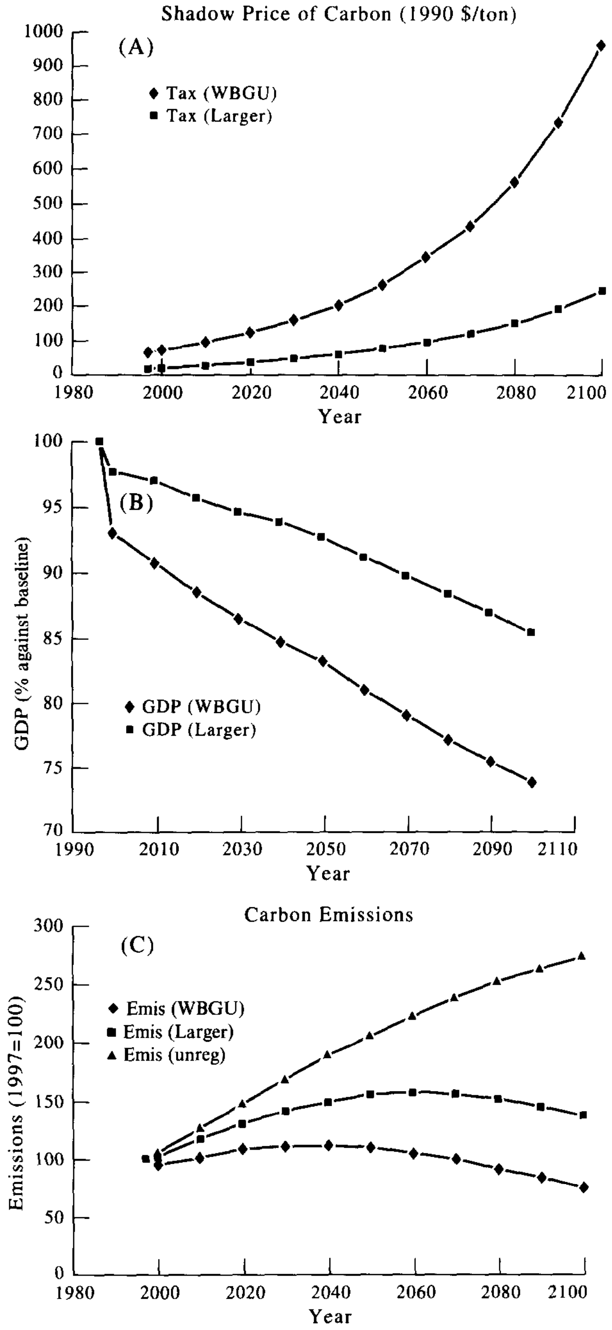

Figure 2. Panel A displays the carbon taxes required to achieve the WBGU and Large windows along the median emissions trajectory. Panel B portrays the associated levels of global GDP for the two windows as a percentage of unrestricted GDP along the median scenario. Panel C compares carbon emissions along the unrestricted trajectory with emissions paths consistent with the two alternative target windows.

two shadow price trajectories for the median case — a lower path for achieving the Larger Window, and a higher path for achieving the WBGU Window. Neither achievement would be possible without cost, though. Table 2 shows that control costs would climb with cumulative emissions, and Figure 2B portrays the corresponding trajectory in global GDP for case 3 in terms of the percentage of unrestricted GDP.

Even though the policy profiles reported in Table 2 are the results of approximations, they are the products of computations that are not unlike the ones that would be made by the actual crafters of mitigation policy who want to mimic what Wigley et al. (1996) (WRE) term 'when' flexibility in an information-constrained decision environment. Solving the scarcity rent problem supports, to a first-order approximation, intertemporal efficiency; and, as shown in Figure 2C, the resulting regulated emissions trajectories have a 'WRE' shape to them. They are also drawn from a aggregate global model; and so they also assume that each ton of carbon removed from the global emissions stream is removed at the least cost site. As such, these results also reflect WRE 'where' flexibility. In short, Table 1 reports reductions in cumulative emissions, maximum temperature change levels, maximum temperature change rates, and minimum control costs.

3. Hedging when neither the target window nor the emissions path is known

The hedging exercise was divided into two steps. In the first, the least (expected) cost strategy was chosen for hedging in 1997 across the five representative emissions trajectories for both target windows. It was assumed, in this first calculation, that:

(1) the likelihood weights assigned to the various scenarios by the Yohe and Wallace (1996) model reflected the subjective view of their relative likelihood in the eyes of the decision-maker who must frame short-term mitigation policy; and
(2) that all the uncertainty about which trajectory most accurately describes the future would be resolved by the year 2020.

The second part of the hedging calculation cast the two resulting least-cost strategies for the two target windows into an environment in which the 1997 decision must be made without knowing which of the two would turn out to be the appropriate window. In hedging across the two targets, though, the decision was assumed to be made with the understanding that target uncertainty would also be resolved by the year 2020.

Table 3 records sufficient data to solve the first hedging problem for both windows. The estimates reported there indicate the discounted cost of meeting the constraint of either window when short-term policies that are consistent with any one specific emissions trajectory of cases

Table 3 Control costs for acting then learning[a]

	Case 1	Case 2	Case 3	Case 4	Case 5	Expected cost[b]
			Achieving the WBGU Window:			
Policy choice:						
as if Case 1	1.2	5.3	23.8	32.5	527.8	107.7
as if Case 2	1.3	5	21	29.1	486.5	98.9
as if Case 3	3.5	6.3	17.4	24.3	407.9	83.8
as if Case 4	3.6	6.4	17.5	24.3	407.7	83.9
as if Case 5	48.4	53.9	46.6	55.7	300.4	95.4
			Achieving the Larger Window			
	Case 1	**Case 2**	**Case 3**	**Case 4**	**Case 5**	**Expected cost[b]**
Policy choice:						
as if Case 1	0	0.5	3	7.1	102.2	20.5
as if Case 2	0.1	0.4	2.6	6.2	90.3	18.1
as if Case 3	0.4	0.6	2.2	5.2	76.1	15.4
as if Case 4	0.5	0.8	2.3	5	72.9	14.8
as if Case 5	6.9	7.6	8.2	8.9	48.9	15.2

[a]Control costs are measured as the present value of deadweight loss through the year 2100 discounted by the Ramsey rate and are denominated in trillions of 1990 dollars.
[b]Expected costs assume the subjective likelihoods of the five trajectories: 0.27, 0.13, 0.23, 0.19, 0.18 for cases (1)–(5), respectively.

(1)–(5) are adjusted over the decade prior to 2020 to accommodate the trajectory that actually comes to pass. For example, the estimate of $17.4 trillion occupying the very middle of the top panel indicates that the present value of control costs incurred in trying to meet the constraints of the smaller WBGU Window by

(1) setting the short-term mitigation policy in 1997 as if the case (3) emissions trajectory were applicable, and
(2) continuing to follow that policy scenario beyond the year 2020 because information available at that time indicated that the case (3) trajectory had actually materialized.

Notice that this cost estimate corresponds exactly to the control cost reported for case (3) in Table 2 for the perfect information 'learn then act' exercise of Section 2. It should be emphasized, therefore, that achieving minimum discounted cost in this more complicated environment is a coincidence of good fortune (being correct in 1997 in anticipating that emissions would actually follow the trajectory of case (3)) and not the product of omniscience. Indeed, all of the estimates noted along the diagonals of both panels conform with the perfect reformation estimates by fortunate coincidence.

The cost of being less fortunate in anticipating the future is also reflected clearly in Table 3. Control cost estimates climb, for the case (3) example, as the actual events deviate from 1997 expectations in any direction. Move to the right of the middle entry in either panel, and the discounted value of control costs climbs because emissions turn out to be moving along a trajectory that is higher than in case (3). Near-term policy that would have been correct for the median case would, in cases like this, be too lax, and so post-2020 emissions policy controls would have to be even more restrictive than otherwise if the target window is to be achieved. Move to the left from the center, and the discounted value of control costs climbs again. In these cases, emissions turn out to be moving along a trajectory that is lower than in case (3), and so near-term policy would have been too restrictive. Excessive costs over the near term would therefore accumulate; and their excess cannot be outweighed by the lower cost of weaker post-2020 policy. Moving down the middle column highlights the discounted cost of setting near-term policy as if a higher emissions profile would develop when, in fact, case (3) will ultimately dictate the future. Near-term policy must, again, be too severe, and discounted costs climb. Finally, moving up the column from the middle picks circumstances in which near-term policy is too lax to handle adequately the emissions of case (3), and so 'catch up' costs incurred past the year 2020 push even their discounted costs up.

The right-hand columns of Table 3 hold the key to the first hedging problem. They reflect the expected cost of choosing short-term policies that are consistent with each of the five alternative emissions trajectories. The calculations assume that the five alternatives have subjective likelihoods of 0.27, 0.23, 0.13, 0.19 and 0.18, respectively. The smallest expected control cost estimate for the smaller WBGU Window is 583.8 trillion when short-term mitigation policy through the year 2020 is crafted as if case (3) would describe the actual long-term future. This minimum expected discounted control cost thus identifies the least-cost hedging strategy for meeting the constraints of the WBGU Window. The smallest expected control cost for the Larger Window is $14.8 trillion when short-term policy is crafted as if case (4) would describe the actual long-term future; and so the least-cost hedging strategy for meeting the constraints of the Larger Window is to prescribe a policy for which the shadow price of carbon will follow the path most consistent with case (4).

Turning now to the second hedging issue, the selection of short-term policy profiles for each window plays a critical role. Indeed, the minimum expected present values of control costs for each target window identified in Table 3 are half of the equation. They are estimates of the expected cost of hedging across alternative emissions trajectories when the policy target window *is known with certainty*. They reflect, more specifically, the minimum expected cost of meeting the constraints of one of the two windows if, in fact, that window turns out to be the correct target. But what if that were not the case? Table 5 provides sufficient information to answer this question. The top panel of Table 5 records, for example, pertinent estimates for the five alternative emissions futures associated with

(1) a short-term policy designed to minimize expected control costs across five probabilistically weighted emissions trajectories with the expectation that the Larger Window would turn out to be appropriate (i.e., setting short-term policy as if the Larger Window were the target along the case (4) emissions trajectory) when, in fact,
(2) information available for incorporation into mitigation policy by the year 2020 supports the notion that the smaller WBGU Window would be more appropriate.

Summing the probabilistically weighted control costs across the five future options produces an estimate for expected, discounted control costs of \$101.7 trillion — an amount far in excess of having planned for the WBGU Window from the beginning (\$83.8 trillion) and even more distant from the expected value of perfect reformation (derived from Table 1 to be \$63.7 trillion).

The bottom panel of Table 5, meanwhile, highlights comparable estimates across the five alternative futures for the opposite situation. In this case,

(1) short-term policy would be crafted to minimize expected control costs across the five probabilistically weighted emissions trajectories assuming that the smaller WBGU Window would turn out to be appropriate (i.e., setting short-term policy as if the WBGU Window were the target along the case (3) emissions trajectory) when, in fact,
(2) information available for incorporation into mitigation policy by the year 2020 would support the conclusion that the Larger Window would be more appropriate.

Expected discounted control costs sum to slightly more than \$15.0 trillion and are slightly higher than the expected discounted control cost of correctly planning for the Larger Window all along (\$14.8 trillion from Table 3). Both are nearly 40% higher than the expected cost of achieving the Larger Window under perfect reformation (\$10.4 trillion from Table 2). Figure 3A and 3B depict the types of adjustment that might be required in this latter case. Panel A plots the trajectory of the shadow price of carbon over time when serious adjustments are required after incorrectly pursuing the WBGU Window along the case (3) trajectory through the year 2020 when, in fact, the appropriate policy would have pursued the Larger Window along the case (2) or (4) scenarios; but it

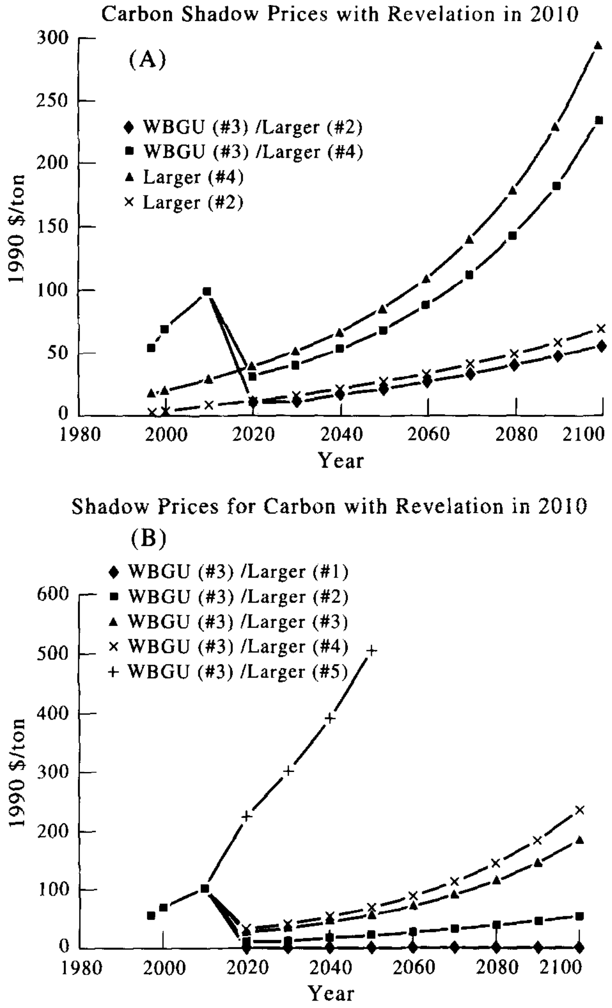

Figure 3. Panel A displays carbon tax trajectories designed to achieve the Large window along one high and one low emissions trajectory after targeting the WBGU Window as if the median trajectory were appropriate through 2010 and compares them with the trajectories that would have been selected if both the actual trajectories and the Large target window had been known in 1997. Panel B portrays comparable adjustments for all of the alternative emissions trajectories for the Large window after seting short-term policy to achieve the WVGU window along the median.

also plots the shadow price of carbon that would be appropriate over time if it had been known in 1997 (a) that either case (2) or (4) would emerge and (b) that the Larger Window would be the appropriate target. Panel B highlights the adjustments required for all five emissions trajectories when the short-term policy incorrectly targets the WBGU Window.

The information collected in Tables 3–5 can now be summarized in terms of two possible hedging options that minimize the expected present value of control costs of targeting one or another window in writing short-term policy. If the WBGU Window were targeted from the start, then the expected present value of control costs would be $83.4 trillion if the WBGU Window turned out to be the correct target and $15.0 trillion if the Larger Window were ultimately selected. By way of contrast, if the Larger Window were targeted in the framing of short-term policy, then the expected present value of control costs in this case would be $101.7 trillion if the WBGU Window turned out to be correct and $15 trillion if the Larger Window were chosen. The corresponding expected present values of assuming perfect information about trajectory and target prior to setting any policy amount to $63,4 trillion and $10.4 trillion for the WBGU and Larger Window targets, respectively. Early information is therefore valuable, but that should have been obvious

Table 4 Cumulative reduction in emissions through 2100 for acting then learning[a]

	Case 1	Case 2	Case 3	Case 4	Case 5
	Achieving the WBGU Window:				
Policy choice:					
as if Case 1	31.9	56.2	62.5	68.9	77.5
as if Case 2	31	55.3	62	68.5	77.5
as if Case 3	29.8	53.7	60.9	67.7	77.7
as if Case 4	29.8	53.7	60.9	67.7	77.7
as if Case 5	33.3	55.1	59.2	67.6	78.8
	Achieving the Larger Window:				
	Case 1	**Case 2**	**Case 3**	**Case 4**	**Case 5**
Policy choice:					
as if Case 1	0	24.5	36.9	50.6	65.6
as if Case 2	9.7	23.4	35.7	49.2	64.8
as if Case 3	11.5	22.3	34.2	47.5	63.8
as if Case 4	12	22.1	33.8	47.08	63.6
as if Case 5	18.3	25.3	33	46.05	61.2

[a]Measured as percentage reductions against unregulated emissions trajectories.

Table 5 Summary statistics for achieving tolerable windows: selected hedging strategies

Case	Emissions reduction	Control cost[c]	Max. temp.	Max. rate
	Hedging with a Larger Window strategy when the WBGU Window obtains:[a]			
1	31.30%	1.3	1.05	0.14
2	55.60%	5	1.01	0.148
3	62.10%	21.8	1	0.151
4	68.60%	30	1.03	0.154
5	77.52%	500.1	0.99	0.157
	Hedging with WBGU Window strategy when the Larger Window obtains:[b]			
Case	**Emissions reduction**	**Control cost[c]**	**Max. temp.**	**Max. rate**
1	16%	3.2	1.9	0.136
2	23.80%	3.6	1.85	0.18
3	32.10%	4.4	1.85	0.18
4	45.33%	5.9	1.82	0.183
5	61.90%	64.5	1.78	0.187

[a]Hedging for the Larger Window involves taking the least-cost hedging strategy through 2020 assuming that the Larger Window will be chosen as the policy target; i.e., imposing a carbon tax as if case (4) will apply in shooting for the Larger Window.
[b]Hedging for the WBGU Window involves taking the least-cost hedging strategy through 2020 assuming that the WBGU Window will be chosen as the policy target; i.e., imposing a carbon tax as if case (3) will apply in shooting for the WGBU Window.
[c]Denominated in trillions of 1990 dollars.

from the start. What is new here is that the expected value of the cost of targeting the WBGU Window in setting near-term policy is smaller than the expected cost of targeting the less restrictive Larger Window as long as the subjective likelihood that the WBGU Window will actually turn out to be the better choice is larger than 11%.

4. Conclusions and context

The results reported here suggest a startling hypothesis: least-cost hedging in anticipation of achieving tolerable windows, defined not only in terms of changes in global mean temperature but also in terms of the rate of change of temperature, can justify significant near-term reductions in carbon emissions. Their specifics suggest, in fact, that near-term emissions might need to be held close to 1990 levels at least through the year 2020. This is perhaps the first time that a cost-based analysis has offered any support for such extreme mitigation; and so the results raise a larger number of questions that need to be addressed before even their qualitative content is to be the least bit credible.

It is, first of all, essential that their robustness be tested by exercising more detailed and disaggregated models. Some of the most immediate questions are essentially economic. How sensitive are the qualitative conclusions to distributional issues, for example? Does burden sharing matter? Can small windows of tolerable change actually be achieved over the near term without complete global participation, especially along trajectories where emissions from developed and/or developing countries run higher than expected? Other pressing questions focus attention on the scientific underpinnings of even the simple model. How sensitive are the conclusions to climate sensitivity? Inasmuch as the rate of change constraint bites early in any trajectory, would they persist if near-term sulfate emissions were tracked carefully? Would the policy prescription persist if sulphate emissions rose significantly somewhere in the world through the year 2020? And what would happen if they fell significantly somewhere else?

The results reported here also offer some insight to the batteries of researchers who are currently engaged in drafting new emissions scenarios for both the IPCC and the next EMF exercise. Most of the action occurs in this analysis along the highest emissions trajectory, but it is not a trajectory that describes a future with minimal potential likelihood. In fact, case (5) has a subjective weight of nearly 20%. Both groups should therefore take care to include a 'baseline' scenario that reflects something like the 75th or 80th percentile emissions path — with the percentiles reflecting a range of uncertainty of futures from models that can offer distributions of outcomes and not a range of disagreement from models that report 'best guess' scenarios.

Perhaps the most pressing questions raised here turn attention to the target windows themselves. This analysis has taken them as given in computing minimum expected control costs, but is that assumption appropriate? Indeed, the expected control cost estimates reported can be interpreted as the expected discounted opportunity cost of being intolerant to rapid change in temperature over the relatively near term. But is that intolerance justified? Everywhere around the world? Toth (1997) notes, for example, that both windows were based to some degree by setting the limit for tolerable sea level rise at 30 cm. This is clearly not a definition of intolerance for developed countries where protection or efficient abandonment of coastal property can be anticipated. It may be a reasonable definition of intolerance for developing countries, but the cost associated with such intolerance is potentially enormous.

Yohe (1997) suggests that the key to assessing intolerance will lie in considering three distinct types of adaptation. The first is adaptation that can make a system less vulnerable to short-term variation in its current environment. Why? Reaction to short-term fluctuation can be a precursor to long-term adaptation strategies; and more robust systems have increased access to the physical and human capital from which long-term adaptation can be built. The second is adaptation over the long term to anticipated long-term change. This is the type that most people envision, and its role is obvious. The third type of adaptation turns this familiar story around by recognizing that systems can switch their points of emphasis and evolve over time so that their sources of sustenance change from resources that will become less abundant to resources that will become relatively more abundant. Adaptation that accelerates this sort of transition by enhancing the ability to make the switch can serve to sustain welfare even as the system itself changes dramatically.

All of the discussion about adaptation notwithstanding, the simple reality is that a future with global climate change and global change policy will not be pareto improving. It need not be totally intolerable for larger numbers of human beings, but adjustments and adaptations will not leave everyone at least as well off as they would have been without climate change. At the breakeven likelihood for the WBGU Window, expected discounted control costs sum to more than $21 trillion — a big number that nearly matches global GDP in 1990. Put another way, 'buying into' the break-even hedge would be the equivalent of commitimg the globe to the purchase of an annuity worth $21 trillion; and such an annuity could, assuming only a 3% discount rate, indefinitely sustain global expenditure in the interest of social welfare, broadly defined, in excess of $650 billion *each and every year*. Do we want to invest the income from that annuity in climate change? Or are there better things to do with that much money?

References

Energy Modeling Forum (EMF-14) (1994) Summary of results. EMF mimeo, Stanford University, CT.

Houghton, J., Callander, B. and Varney, S. (Eds.) (IPCC, 1992) *Climate Change 1992: The Supplementary Report to the IPCC Scientific Assessment.* Cambridge University Press, Cambridge, UK.

Koopmans, T. (1967) Objectives, constraints and outcomes in optimal growth models. *Econometrica* **35**, 1–15.

Manne, A. (1996) Hedging strategies for global carbon dioxide abatement: a summary of poll results. EMF 14 Subgroup; Analysis of decisions under uncertainty, Energy Modeling Forum Working Paper 14.1, February.

Nordhaus, W. and Yohe, G. (1983) Future carbon dioxide emissions from fossil fuels. In *Climatic Change*, ed. National Research Council. National Academy Press, Washington, DC.

Rotmans, J. and Dowlatabadi, H. (1996) *Integrated Assessment of Climate Change: Evaluation of Models and Other Methods* (with E. Parson, contributing author). RIVM mimeo.

Toth, F. (1997) The tolerable windows approach to integrated assessments. *IPCC Asia–Pacific Workshop on Integrated Assessment Models*, Tokyo, Japan, 10–12 March. Mimeo, Potsdam-Institut für Klimafolgenforschung.

Weyant, J. (1997) Insights from integrated assessment. *IPCC Asia–Pacific Workshop on Integrated Assessment Models*, Tokyo, Japan, 10–12 March.

WBGU (German Advisory Board on Global Change) (1995) *Scenario for the Derivation of Global CO_2 Reduction Targets and Implementation Strategies.* Bremerhaven.

Wigley, T.Richels, R.Edmonds, J.$_2$ concentrations (1996) Economics and environmental choices in the stabilization of atmospheric CO. *Nature* **379**, 240–243.

Yobe, G. (1996) Exercises in hedging against extreme consequences of global change and the expected value of information. *Global Environmental Change* **6**, 87–101.

Yohe, G. (1997) Assessing the role of adaptation in evaluating vulnerability to climate change. *International Workshop on Climate Change: Thresholds and Response Functions*, Potsdam-Institut für Klimafolgenforschung, Potsdam, Germany, 23–24 January.

Yohe, G. and Wallace, R. (1996) Near term mitigation policy for global change under uncertainty: minimizing the expected cost of meeting unknown concentration thresholds. *Environmental Modeling and Assessment* **1**, 47–57.