When We Don't Know the Costs or the Benefits:

Adaptive Strategies for Abating Climate Change

Robert J. Lempert[1], Michael E. Schlesinger[2] and Steve C. Bankes[1]

[1] RAND
1700 Main St.
Santa Monica, CA 90407

[2] Department of Atmospheric Sciences
University of Illinois, Urbana-Champaign
105 South Gregory Ave.
Urbana, IL 61801

Climatic Change, 33, 235-274

ABSTRACT

Most quantitative studies of climate-change policy attempt to predict the greenhouse-gas reduction plan that will have the optimum balance of long-term costs and benefits. We find that the large uncertainties associated with the climate-change problem can make the policy prescriptions of this traditional approach unreliable. In this study, we construct a large uncertainty space that includes the possibility of large and/or abrupt climate changes and/or of technology breakthroughs that radically reduce projected abatement costs. We use computational experiments on a linked system of climate and economic models to compare the performance of a simple adaptive strategy – one that can make midcourse corrections based on observations of the climate and economic systems – and two commonly advocated 'best-estimate' policies based on different expectations about the long-term consequences of climate change. We find that the 'Do-a-Little' and 'Emissions-Stabilization' best-estimate policies perform well in the respective regions of the uncertainty space where their estimates are valid, but can fail severely in those regions where their estimates are wrong. In contrast, the adaptive strategy can make midcourse corrections and avoid significant errors. While its success is no surprise, the adaptive-strategy approach provides an analytic framework to examine important policy and research issues that will likely arise as society adapts to climate change, and which cannot be easily addressed in studies using best-estimate approaches.

1. INTRODUCTION

To celebrate the World's Colombian Exposition in Chicago in 1893, the American Press Association asked seventy-four noted commentators from many fields to predict what American life would be like in the 1990's (Walter, 1992). Today, the range in accuracy of these essays seems striking. Some writers projected key trends and produced visions accurate in concept, though rarely in detail, while other essays amusingly extrapolated trends that have since radically changed. For instance, one essay envisions that by the 1990's most businesses will communicate by means of electric transmissions, while another envisions that only the poor will use doctors.[1] In this group of 74 experts there were some who were reasonable seers. However, in 1893, it was not possible to know which ones they were.

The lack of accurate foresight often poses severe challenges for policy-makers. In particular, the need to look far into the future greatly complicates our attempts to estimate the costs and benefits of alternative policies for addressing the threat of climate change. The time scales inherent in the climate and economic systems mean that our decisions today can still have implications decades from now. Currently available, decision-analysis tools generally suggest an optimum policy based on a best-estimate prediction of the future. Many of these studies explore the massive uncertainty associated with climate change by treating these best estimates as probability distributions with known mean and variance, and then finding the optimum policy as a function of these distributions (Nordhaus, 1994a; Peck and Teisberg, 1993; Manne and Richels, 1992).

The time scales of the climate-change problem require that analysts run the computer models used in such cost-benefit approaches a century or more into the future so that the boundary conditions do not affect the answers. But in so doing, recommendations for the optimum near-term policy will in general depend on the predictions we make about what the climate and society will be like at the middle or end of the next century – for instance, how likely it is that the climate will change, that society will be very sensitive to these changes, or that technologies will become available to significantly reduce greenhouse-gas emissions at little cost? Despite the power of the traditional methods, we should feel uneasy if their policy recommendations depend too strongly on particular best estimates of a distant future.

[1] Editor and author John Habberton made this erroneous prediction about access to health care. He argued that "medicine will be practiced at police stations and among outcasts, for respectable people will have resolved that illness not caused by accident is disgracefully criminal." John Wanamaker, Postmaster General under President Benjamin Harrison, made the surprisingly accurate prediction about future communications.

In this paper we explore an alternative approach. We start with the assumption that, while we can envision many plausible futures relevant to today's choices about climate-change policy, we have no way of knowing which of those predictions will turn out to be correct. We then show that a simple adaptive strategy, designed to be robust across many plausible futures, performs better on average than policies optimized for particular best-estimates of the future, unless we are virtually certain that one best-estimate is correct. Thus, the adaptive-strategies framework may help policy-makers make reasonable and defensible choices about near-term climate-change policy without requiring accurate or widely accepted predictions of the future.

This result also suggests that society might usefully recast its view of the climate change problem. Currently, the political debate over climate-change policy focuses on the targets and timetables for optimum level of near-term reductions in greenhouse-gas emissions that society should or should not set. The research community views their task as improving the accuracy of the predictions of the future which will provide policy-makers with better estimates of the optimum level of near-term emissions reductions (Global Change Research Program, 1995). The approach here suggests climate change be viewed as a problem of preparing for unpredictable contingencies. Society may or may not need to implement massive reductions in greenhouse-gas emissions in the next few decades. The problems for the present include developing better options for massive reductions than those currently available and determining what observations ought to trigger their implementation. We believe that this study presents a first step towards an analytic framework for addressing such questions.

2. Exploratory Modeling and Adaptive Strategies

The traditional decision-analytic approach to the climate-change problem has produced some powerful insights. It has stressed the importance of balancing costs and benefits in any climate-change policy; emphasized the value of more research on climate change; and suggested that given our current understanding, any near-term reductions in greenhouse-gas emissions ought to be moderate. These insights have injected important balance into the public climate-change debate (Passell, 1989).

Nonetheless, these approaches have significant shortcomings that limit their ability to address the full range of policy issues raised by the climate-change problem. As we will describe below, these are not so much conceptual problems as issues of the allocation of computational resources.

In principle, the traditional methods can address each of the issues we raise. In practice, the computational demands of calculating the optimum policy for even a single best estimate forces the analyst to neglect aspects of the problem central to some important questions of climate-change policy. We argue here that the traditional approach should be complemented with a different set of analytic tools to give a more complete view of the policy problem and the options available for its solution.

The first problem is that the traditional approach has difficulty grappling with the immense range of uncertainty inherent in our understanding of climate change. For instance, many such methods confine themselves to power-law damage functions (usually quadratic, linear, or cubic) to facilitate computations. While power-law functions are generally good approximations for small excursions around point estimates (being merely the terms of the Taylor expansion), they can be very poor approximations for large excursions over unknown functions, particularly ones that might exhibit non-linear or abrupt changes.

As an example, Figure 1 shows a range of damage functions that Nordhaus (1994a) considers in his landmark study of climate-change policy. These functions are the result of an exhaustive examination of the range of uncertainties, as well as the addition of an extreme case used to explore the consequences of catastrophic damage. While these damage functions span a large range of values, they all have the same shape (except for the identically zero first-quintile estimate). The first and second derivatives are all monotonically increasing, so that each function shows damages due to climate change that are very much smaller in the first half of the 21st century than in the second half. Traditional approaches usually confine themselves to damage functions of this shape to facilitate the computation of an optimum policy.

Nordhaus uses the damage functions in Figure 1 to argue that society should reduce emissions very little over the next few decades. There is, however, no reason to believe that the damages due to climate change – which implicitly includes society's ability to adapt to any changes that occur – must be a convex function. In this paper, we will show that under certain circumstances, non-convex damage functions can have a significant impact on the optimum rate of near-term abatement. Thus, the computational difficulties of the traditional approach discourage the treatment of plausible futures – non-convex damage functions – that directly affects one of the key issues addressed by these studies – the desirability of near-term abatement.

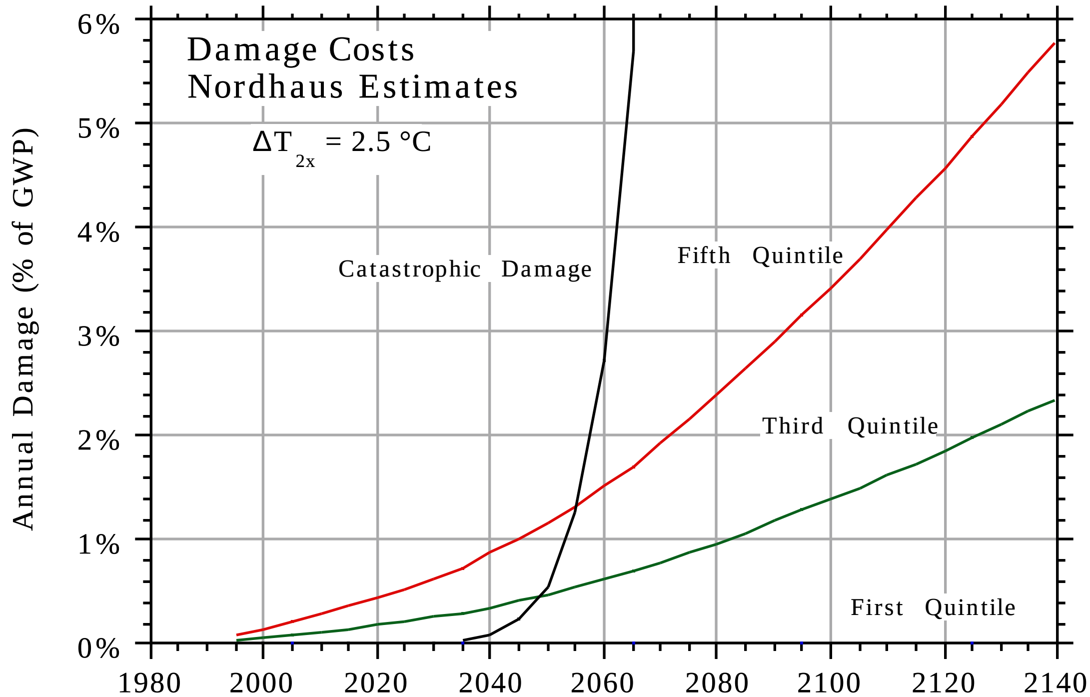

Figure 1: Estimates of annual damage costs from Nordhaus (1994a) using the damage function d = q1 (ΔT/3)2 with q1 = 1.3%, 0%, and 3.2% of Gross World Product, the best estimate, first quintile, and fifth quintile values, respectively, of the damage parameter q1 (see Nordhaus' Table 7.1), using the ΔT2x = 2.5 °C, 'Conservation Only' temperature trajectory from Figure 4a. Also shown are estimates of annual damages using Nordhaus' catastrophic damage function, d = 0.27(ΔT/2.5)12, from Chapter 6.

The second problem is that best-estimate approaches often treat uncertainty in one place by assuming certainty someplace else. For instance, these approaches assume that when we cannot predict some value or function, such as the damages due to climate change, we can use a well-defined probability distribution to represent the range of possibilities. The traditional approach was developed for problems where the distribution of outcomes is well-known, even when the outcome themselves are unpredictable. For instance, an aircraft designer cannot predict when an individual engine might fail, but can draw on society's extensive experience with aircraft engines to accurately estimate the probability that an aircraft's engines will fail in flight. For many policy problems, such objective probabilities (based on repeated observations of a given system) are not available, but we can often use subjective probabilities that represent our best estimate of the probability distribution of outcomes. Often subjective probabilities are developed from surveys of the relevant scientific experts (Morgan and Herion, 1990).

In the climate-change problem, however, the scientific community may know as little about the probability that some outcome will occur as it does about the outcomes themselves. In this paper we will show that the choice among best-estimate climate-change policies depends strongly on the subjective probabilities assigned to different plausible futures. Policy-makers should be wary of policy recommendations that depend strongly on subjective probabilities which may themselves be highly uncertain. In addition, any policy response to climate change will likely engage the interests of many different stakeholders. We can expect that each group will hold different subjective probabilities, which likely will reflect their particular institutional interests.

At some point, of course, policy-makers must make their own guesses about the future when choosing a climate-change policy. Nonetheless, both because experts have little solid knowledge sufficient to assign likelihoods to different plausible futures, and because competing stakeholders will cling to different best-estimates, an approach that generates optimal policies which depend strongly on the subjective probabilities may have limited influence in actual policy debates. Policy analysts may be most helpful when suggesting sub-optimal approaches that promise to be reasonably effective given many different visions of the future.

The third problem is that the traditional approach must often neglect important feedbacks among different elements of the climate and economic systems. The traditional approach is well-designed to treat the important set of feedbacks that produce optimal growth in competitive market economies in the presence of environmental externalities. Among the important insights gleaned from such studies is that carbon taxes may be the most efficient means for addressing climate-change externalities, since the optimal carbon tax is less sensitive to the resolution of uncertainties than is the optimal rate of greenhouse-gas abatement (Nordhaus, 1994a).

There are many other feedbacks, however, that may also have significant effects on our choice of climate-change policies, but which are difficult to treat in the traditional approach. Such feedbacks may make optimization approaches analytically intractable. Also, in many cases there is no model that the analyst would feel comfortable presenting as a best estimate of the phenomena. In this study, we will examine the important feedback of society's ability to gain information about the seriousness of climate change and the costliness of potential responses, this by observing the evolution of the climate and linked economic systems over the coming decades.

In this paper we demonstrate a new set of quantitative tools that offer the potential to address important climate-change policy questions which are difficult to treat with the traditional approaches. Our approach is based on two concepts that we call adaptive strategies and exploratory modeling.

Exploratory modeling treats problems having massive uncertainty by conducting a large number of computer simulation-experiments on many plausible formulations of the problem, rather than using computer resources to increase the resolution of a single best-estimate model (Bankes and Gillogly, submitted; Bankes and Gillogly, 1994a; Bankes and Gillogly, 1994b, Bankes, 1994; Bankes, 1993). Exploratory modeling calls for explicitly modeling a large number of potentially salient uncertainties, forming from these uncertainties a space or ensemble of possible computational experiments, and then devising sampling strategies for this ensemble to address particular questions.

As we apply it here, exploratory modeling seeks to find strategies that are robust to the uncertainty society faces over climate change – that is, seeks to find strategies that perform reasonably well over broad ranges of plausible futures – rather than find the policy which is optimum for some particular set of subjective probabilities. Exploratory modeling also allows freer use of non-linear functional forms and complicated feedbacks than would be the case in most optimization approaches. In general, exploratory modeling represents a reallocation of computational resources away from finding the optimum policy for a small number of cases to comparing the performance of alternative strategies over a large number of plausible futures.[2]

Adaptive-decision strategies focus on modeling environmental policies where decision-makers can make midcourse corrections based on observations of the relevant environmental and economic systems. The adaptive strategy is a generalization of the sequential-strategy approach in our previous work (Lempert, Schlesinger, and Hammitt, 1994 [henceforth LSH]; Hammitt, Lempert, and Schlesinger, 1992 [henceforth HLS]) and the 'Learn, then act' framework of Manne and Richels (1992) and Nordhaus (1994a). Sequential strategies can make midcourse corrections at some fixed time in the future based on some improvement in information that occurs independently of what is happening in the climate system. The sequential framework can provide important results. For instance, in LSH and HLS we found that the long-term costs of responding to climate change, whether or not there are abrupt changes, is relatively insensitive to whether we choose aggressive or moderate abatement over the next ten years. However, the sequential-strategy approach is limited because it cannot treat the interplay

[2] While an exploratory-modeling analysis often requires significant computing resources, the calculations can be conveniently configured as a parallel process which can run simultaneously on a distributed network of computer workstations. This study required 38,884 computational experiments, each consisting of a 156-year scenario (1995-2150), which we made by exploiting otherwise idle processor time on a network of 5 workstations (Sun Sparc 2’s, 10’s, and 20’s) over a period of 5 days, thereby obtaining supercomputer levels of processing power without the use or cost of a supercomputer.

between society's rate of learning and the rate of change in the climate system.

Nordhaus' discussion of the catastrophic damage function in Figure 1 provides a particularly graphic example of the importance of this information feedback. Nordhaus argues that such a damage function calls for only moderate emission reductions in the near-term, with rapidly increasing reductions in the far-term to keep society away from the threshold. "It is crucial," Nordhaus writes, "for this analysis that policy-makers are aware of the threshold." But what if they are not? If we assign a probability p that society will learn about some catastrophic threshold before crossing it, the optimum near-term emissions-abatement policy will range from no reductions to draconian reductions as a function of p. This is not a very useful result since we have no way of knowing, much less agreeing as a society, upon the value of p. However, this analysis also overstates the case. As society approached some catastrophic threshold, it would likely observe hints of its approach. The question becomes, could society recognize and act on these hints in time? More generally, might society reach some consensus on near-term actions to increase the chances it could recognize and respond to avoid serious consequences of climate change? Such questions – the interplay between learning rates and response rates and the policy actions that can change these rates – emerge as central issues in the adaptive-strategies framework. They are not impossible to treat in current approaches,[3] but they are computationally awkward in methods designed to produce optimum policies for well-understood systems.[4]

This paper reports on an initial attempt to apply an exploratory-modeling/adaptive-strategies analysis to the climate-change problem. The current policy debate is largely divided into those who favor the emissions-stabilization goals laid out in the 1992 Rio Treaty and those who advocate minimal near-term emissions reductions. In this work we construct a large uncertainty space representing many plausible futures, that encompasses order-of-magnitude differences in the climate sensitivity, in the damages due to climate change, and in the possibilities for innovation to radically lower abatement costs. We show that the choice between the two commonly advocated 'best-estimate' policies – 'Do-a-Little' and 'Emissions-Stabilization' – largely depends on the subjective probabilities one places on different long-

[3] Kolstad (1994, 1993) addressing learning and irreversibilities within an optimization approach.

[4] Nordhaus (1994a), Manne and Richels (1992), and Peck and Teisberg (1993) also use the sequential strategy approach to estimate the value of information about various climate variables. Value of information can also be easily treated in the adaptive-strategy approach. We do not do so here, because the value of information to an adaptive strategy depends on the level of noise and other fluctuations in the system. This is a subject of our ongoing research.

term predictions of the future. We are unlikely to predict these futures correctly, and the cost of choosing the wrong best-estimate policy can be very high. Thus, the choice between 'Do-a-Little' and 'Emissions-Stabilization' is not very appealing, nor does it appear to be a good framework in which to resolve the debate over climate-change policy.

In this work we also compare the performance of a simple adaptive strategy, with moderate near-term abatement,[5] to the performance of the two best-estimate policies. We find that the adaptive strategy, because it can make midcourse corrections and avoid significant errors, performs better on average than the best-estimate policies for virtually the entire set of expectations about the future. These results suggest that an adaptive-strategy framework may make it easier for policy-makers to agree upon and craft a robust climate-change policy, even in the absence of any accurate and/or widely shared consensus about the future. The adaptive-strategy approach also casts climate change as a problem of preparing for unpredictable contingencies. This view emphasizes questions – what near-term policies can improve the available options for implementing such aggressive abatement if we learn it's needed? what observations ought to trigger aggressive abatement of greenhouse gases? – that may be easier to treat in the exploratory-modeling/adaptive-strategies framework than with the traditional best-estimate methods.

3. Modeling the Climate/Economic System

In this work we consider the linked system of simple climate and economic models shown in Figure 2. Emissions of greenhouse gases determine their atmospheric concentrations, which in turn determine the change in global-mean surface temperature. These temperature trajectories determine the trajectory of damage costs, while the emissions trajectories generate a trajectory of abatement costs. As described below, we expand on the models used in our previous work (LSH, HLS) to treat these systems.

We model our imperfect predictions about the future evolution of the climate and the economy by three uncertainties that appear particularly important in shaping society's climate-change policy choices (Nordhaus, 1994a; Lave and Dowlatabadi, 1993). These are: (i) the sensitivity of global-mean surface temperature $\Delta T(t)$ to increases in greenhouse gases, which is treated in the climate model; (ii) damages resulting from an increase in global-mean surface temperature, which is treated in the damage model;

---

[5] This paper does not consider the 'optimum' level of near-term abatement for adaptive decision-strategies for climate change. It is not meaningful to address this issue without considering the effects of noise and other system fluctuations on the ability of the adaptive strategy to make unambiguous observations (see Section 5). This is a subject of our ongoing research.

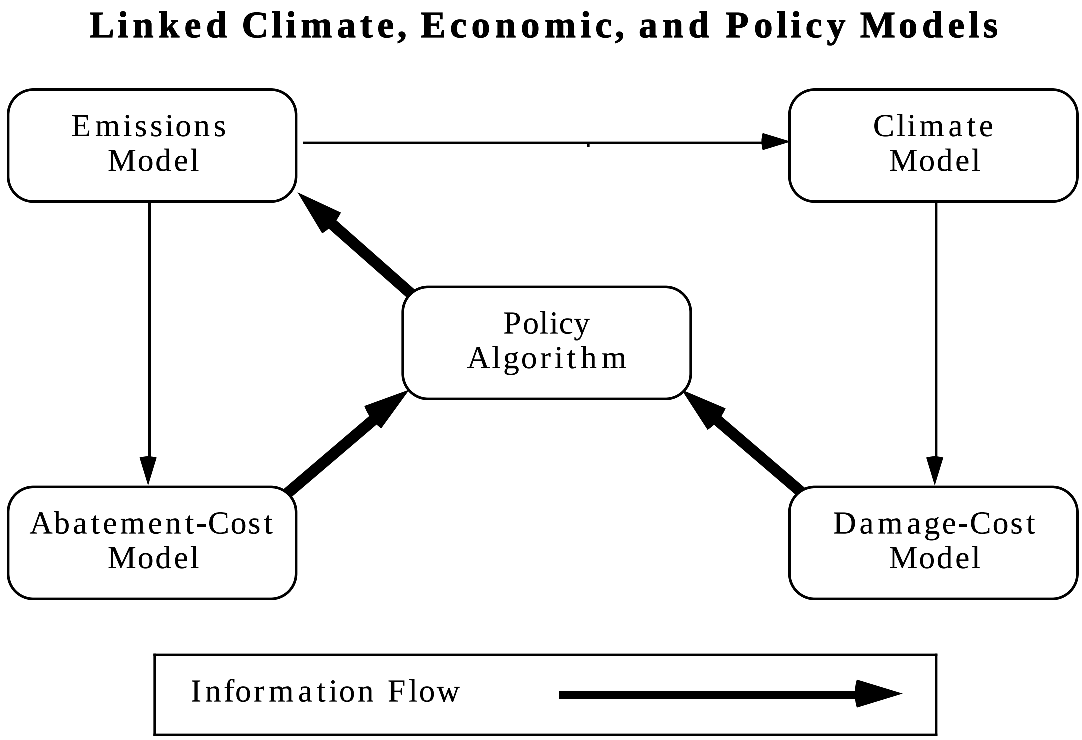

Figure 2: System of linked models examined in this study.

and (iii) the ability of innovation to reduce the cost of abating greenhouse-gas emissions, which is treated in the abatement-cost model. We consider order-of-magnitude uncertainty in each of these systems.

A key aspect of our analysis is that we consider adaptive-policy algorithms that can make midcourse corrections based on observations of the climate and economic systems.[6] We explicitly consider the decision-maker as part of the linked system of models in order to examine the potential impact of information feedbacks. For instance, we can explore the conditions under which rapidly rising damages due to climate change provide warning for an effective policy response. We can thus explore the potential advantages of adaptive strategies compared to policies that do not use observations to make their decisions.

This section provides detailed descriptions of the emissions, climate, abatement-cost, and damage-cost models shown in Figure 2 and our treatment of the uncertainties. In Section 4, we compare the performance of

[6] ICAM-2, the Integrated Climate-Assessment Model developed at Carnegie Mellon University, also treats policies which respond to observations (Parson, 1994).

different policy algorithms.

3.1. Emissions Model

In this work we explicitly consider abatement policies that affect the anthropogenic emissions of carbon dioxide and methane, two of the most important greenhouse gases. (As described below, we also consider the impact of projected emissions of nitrous oxide, CFCs, and CFC-substitutes on global-mean surface temperature, but do not consider policies that change the emissions of these gases.) As in our previous work (LSH, HLS), we will use a heuristic model, based on the link between energy-use and greenhouse-gas release, to simulate the effect of stylized policy choices on emissions of carbon dioxide and methane and the abatement costs associated these emissions trajectories.

We model the emissions of carbon dioxide and methane with the equation,

$$F_{\sigma}(t) = B_{\sigma}(t) I(t) E(t) \quad , \tag{1}$$

where $\sigma = CO_2, CH_4$ for carbon dioxide and methane, respectively. The $B_{CO_2}(t)$ and $B_{CH_4}(t)$ are base-case global-aggregate emissions trajectories extracted from the IPCC/90 Scenario A (Houghton et al., 1990), as described in LSH. I(t) and E(t) are functions representing the diffusion of 'conservation' technologies, which decrease the energy intensity (energy used per unit economic activity), and of 'fuel switching' technologies, which decrease the emissions intensity (emissions per unit energy generated) of society's energy-related capital stock, respectively. We are particularly interested in characterizing the rate at which these policies can be implemented, so that we constrain I(t) and E(t) to follow the logistic curves that are characteristic of technology diffusion (Mansfield, 1961; Fisher and Pry, 1971; Hafele, 1981).

Note that we use the same I(t) and E(t) in Eq. (1) for both $F_{CO2}(t)$ and $F_{CH4}(t)$, which assumes that methane emissions are reduced together with carbon-dioxide emissions at no additional cost. This simplification restricts our ability to compare policies that reduce methane and carbon dioxide independently, but should not greatly affect any of the conclusions of this study.

The energy-conservation function I(t) represents a set of low-cost, quickly implemented policies. We assume that the conservation policies are initiated in the year 1995, and that the energy intensity thereafter falls 30% over 20 years. Thus I(t) is given by

$$I(t) = \begin{cases} 1 & \text{for } t < 1995 \\ \Psi(t;1995,10,0.3) & \text{for } t \geq 1995 \end{cases} , \tag{2}$$

where $\psi(t;\, t_o, r, \Lambda)$ is the logistic function

$$\psi(t;\, t_o, r, \Lambda) = (1 - \Lambda) + \frac{1}{1-\varepsilon} \frac{\Lambda}{1 + \exp[\rho_r(t - t_o - r)]} , \tag{3}$$

where r = 10 is the transition half-life and L = 0.3 is the amount of low-cost conservation available. To satisfy I(1995) = 1, we set $\rho_r = -(1/r)\ln[\varepsilon/(1-\varepsilon)]$ and choose e = 0.01 to obtain reasonable slopes at $t_o + r$ (Hafele, 1981). This amount of conservation is an intermediate estimate of what can be accomplished at low cost in the United States (NAS, 1985; OTA, 1991; Fickett et al., 1990). As such, it is probably an optimistic estimate for low-cost conservation available worldwide.

The fuel switching function E(t) represents a set of high-cost, slowly implemented emission-reduction policies. Their costs and effectiveness are simulated by assuming that all emissions are produced by long-lived capital equipment using either emitting (fossil fuel) or non-emitting (e.g., nuclear, solar, biomass) technologies. For both technologies the construction period is ten years and the maximum operating period is thirty years. Similarly to our previous work (LSH, HLS), we assume that in 1995 a decision is made that fuel-switching will begin in 2005 (after ten years of construction) with a transition half-life of R_1, and that the fuel-switching may subsequently be adjusted to a transition half-life of R_2 in the year T_{12}. Thus, the emissions intensity E(t) is given by

$$E(t) = \begin{cases} 1 & \text{for } t < 2005 \\ \Psi(t;2005,R_1,1) & \text{for } 2005 \leq t \leq T_{12} \\ \Psi(t;\tau_{12},R_2,1) & \text{for } t \geq T_{12} \end{cases} , \tag{4}$$

where the logistic $\Psi(t, t_o r, \Lambda)$ is given in Eq. (3) with $\rho_r = -(1/r)\ln[\varepsilon/(1-\varepsilon)]$ and e = 0.01 as above. To ensure continuity at $E(t = T_{12})$, we choose $\tau_{12} = T_{12} - (R_2/R_1)(T_{12} - 1995)$.

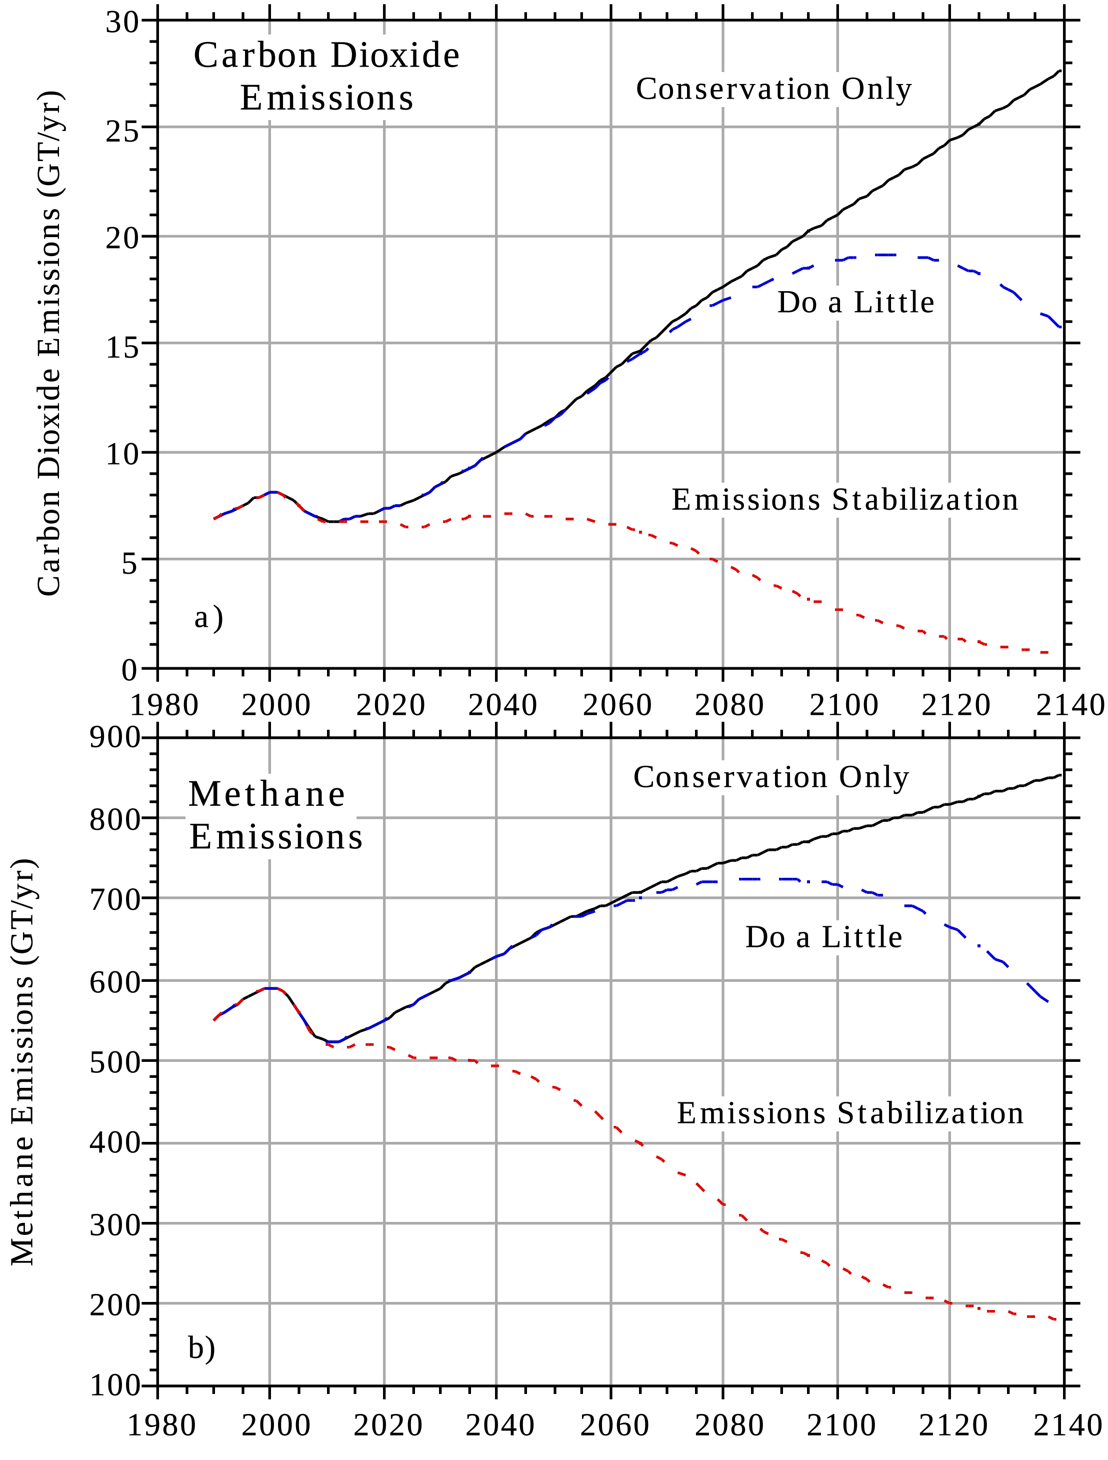

Figure 3: Time evolution of a) carbon-dioxide emissions and b) methane emissions, for the emissions policies 'Conservation Only' (∞/ ∞/na), 'Do a Little' (∞/100/2035), and 'Emissions Stabilization' (30/100/2015).

Figures 3a and 3b show the emissions trajectories for carbon dioxide and methane, respectively, for three emissions policies of particular interest to this study. The 'Conservation-Only' case, defined by (R_1, R_2, T_{12}) = (∞/ ∞/na), is the trajectory relative to which all our abatement costs are quoted. (Note: the ‘na’ signifies that the value of T_{12} is not applicable for emissions policies where $R_1 = R_2$). As we will discuss in more detail in Section 4, the case (R_1, R_2, T_{12}) = (∞/100/2035) is a ‘Do-A-Little’ abatement policy, similar to the optimum policies reported by most economic cost/benefit analyses, and the case (R_1, R_2, T_{12}) = (30/100/2015) is close to the ‘Emissions-Stabilization’ abatement policies called for in many nations’ commitments under the 1992 Rio Treaty.

3.2. Climate Model

The climate model calculates the change in global-mean surface temperature due to the emission trajectories for carbon dioxide and methane. Our first step in deriving the global-mean temperature is to calculate the atmospheric concentrations of these greenhouse gases. The former is derived from the emissions $F_{CO_2}(t)$ with the linear-impulse response function of Maier-Reimer and Hasselmann (1987) [LSH Eq. (2)], whose parameters are obtained by a fit to a coupled ocean-atmosphere model, which includes the effects of the primary ocean sinks, but not of the biota or any modification of the carbon cycle accompanying climate change. We calculate the atmospheric concentration of methane due to $F_{CH_4}(t)$ and the natural emissions $F^{nat}(t) = 0.165 \text{ Gt } CH_4 / \text{yr}$ (Rotmans, 1990) with a first-order difference equation [LSH Eq. (5)], where the decay rate is chosen to reproduce the Scenario-A projections for methane concentrations in IPCC/90.

The radiative forcings due to these concentrations of carbon dioxide and methane are given by the expressions in LSH, Equations (1), (3), and (4), drawn from IPCC/90. As in LSH, we also treat the forcing due to nitrous oxide, CFCs, and CFC-substitutes using the IPCC/90 Scenario-C radiative forcings for these gases for $t < 2100$, and holding the forcing from these gases constant at the 2100 level for $t > 2100$.

We calculate the change in global-mean surface temperature, $\Delta T(t)$, for each radiative-forcing scenario with the energy-balance climate/upwelling-diffusion ocean model of Schlesinger and Jiang (1991). This model calculates the temperature trajectory as a function of the climate sensitivity, ΔT_{2x}, the equilibrium value of global-mean surface temperature change resulting from an atmospheric concentration of carbon dioxide twice its pre-industrial level. As in LSH and HLS, we aggregate into this one parameter all the uncertainties about the response of the climate system to anthropogenic emissions of greenhouse gases. We consider four values, ΔT_{2x} = 0.5, 1.5, 2.5, and 4.5°C. The values 1.5 and 4.5 are the low and high estimates from IPCC/90, the value 2.5 is the IPCC/90 best-estimate, and the value 0.5°C is one of the lowest found in the literature and is due to Lindzen (1990). Figure 4 shows the temperature trajectories due to the 'Conservation-Only' (Figure 4a) and 'Emissions Stabilization' (Figure 4b) abatement policies shown in Figure 3 for each of the 4 values of the climate sensitivity.

3.3. Abatement-Costs Model

The trajectory of abatement costs is determined by our calculated emission trajectories of carbon dioxide and methane. The incremental annual cost of fuel-switching in billions of 1990 dollars is

$$K(t) = C(t) + O(t) \quad , \tag{5}$$

where

$$C(t) = \sum_{j=0}^{2} \kappa_j(t) \sum_{i=0}^{-9} n_{j,i}(t) \tag{6}$$

is the total cost in year t of having under construction $n_{j,\,i}$ plants scheduled for completion in year t – i having emitting (j = 0), low-cost non-emitting (j = 1) and high-cost non-emitting (j = 2) equipment, and

$$O(t) = \sum_{j=0}^{2} \sigma_j(t) \sum_{i=1}^{30} n_{j,i}(t) \tag{7}$$

is the total operating cost in year t of $n_{j,\,i}$ plants having age i in year t. The construction and operating costs for each type of equipment, $\kappa_j(t)$ and $\sigma_j(t)$, are one of the key uncertainties treated in this work and are discussed below.

As in LSH, the values of the $n_{j,\,-9}(t)$ for t ≥ 1990 are determined by I(t) and E(t) from a demand constraint,[7]

$$\sum_{i=-9}^{20} \sum_{j=0}^{2} n_{j,i}(t) = Q_o B_{CO_2}(t + 10) I(t + 10) \quad , \tag{8}$$

[7] We have continued to use the year 1990 to initialize the $n_{j,-9}(t)$, as in HLS. This choice has negligible impact on our results.

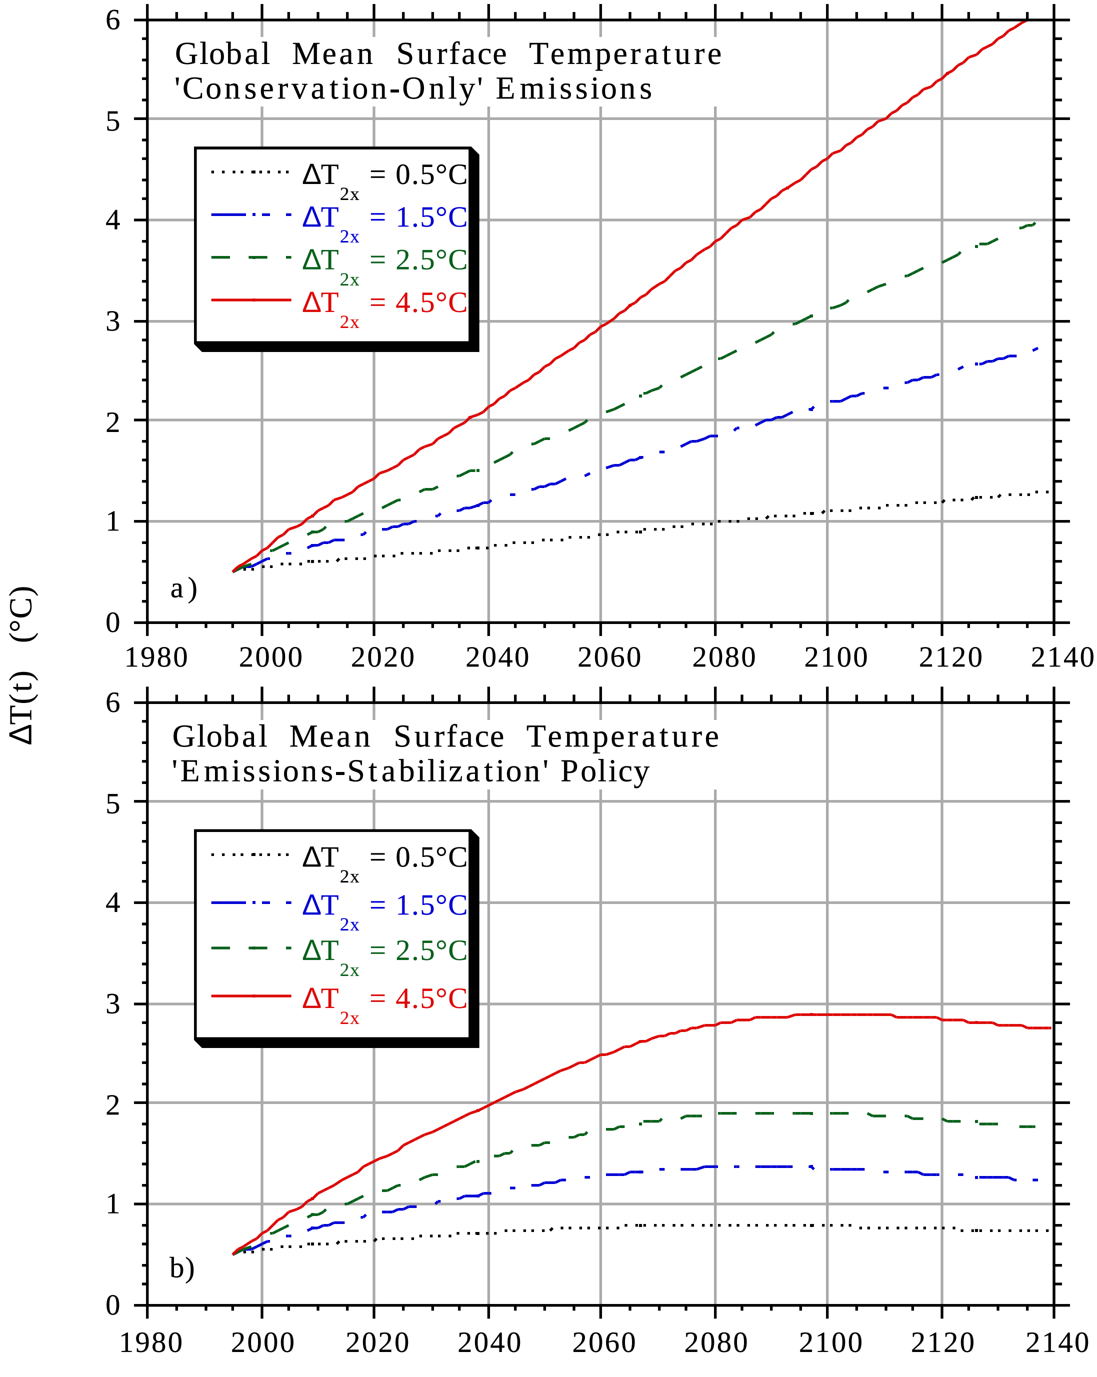

Figure 4: Time evolution of increase in global-mean surface temperature for a) 'Conservation Only' (∞/ ∞ /na) and b) 'Emissions Stabilization' (30/100/2015) policies, for the climate sensitivities ΔT2x = 0.5, 1.5, 2.5, and 4.5°C.

an emission constraint,

$$\sum_{i=-9}^{20} n_{1,i}(t) = Q_o B_{CO_2}(t+10) I(t+10) E(t+10) \quad , \tag{9}$$

and a constraint on the availability of low-cost non-emitting equipment,

$$\sum_{i=-9}^{20} n_{2,i}(t) \le \frac{1}{2} Q_o B_{CO_2}(t+10) I(t+10) \quad , \tag{10}$$

where Q_o is the magnitude of the energy-using sector in 1990, determined as 1990 world CO_2 emissions (6.8 GtC) divided by the 1990 ratio of industrial CO_2 emissions to commercial energy use, 7.2 ton carbon per 105 kWh. For $-8 \le i \le 30$ and all j, $n_{j,i}(t+1) = n_{j,i-1}(t)$. For $t < 1990$ and all i, $n_{1,i}(t) = n_{2,i}(t) = 0$ and $n_{0,i}(t)$ are chosen to reproduce the pre-1990 CO_2 fluxes. Unless Eq. (10) binds, $n_{2,-9}(t) = 0$. If $n_{0,-9}(t) = 0$ is insufficient to satisfy Eq. (10), emitting equipment of ages $i \le 20$ is scheduled for early retirement, oldest first.

The coefficients for the construction and operating costs are given, respectively, by

$$\kappa_j(t) = \kappa_o + (\kappa_j - \kappa_o)(1-d)^{t-t_{tech}} \quad \text{for } t \ge t_{tech},\ j = 1,2 \tag{11a}$$

$$\sigma_j(t) = \sigma_o + (\sigma_j - \sigma_o)(1-d)^{t-t_{tech}} \quad \text{for } t \ge t_{tech},\ j = 1,2 \quad , \tag{11b}$$

where the values of κ_j and σ_j are shown in Table I. These coefficients for the construction and operating costs of low-cost and high-cost non-emitting equipment are equivalent to the \$50 and \$200 per ton of carbon avoided (Manne and Richels, 1991; Nordhaus, 1991).

Predictions of the cost of new technology systems decades into the future are highly uncertain. Thus we include in Eqs. (11a) and (11b) the effects of innovation which reduces the incremental cost of abating carbon dioxide and methane emissions by d percent per year starting in the year t_{tech}. We consider three values of this innovation parameter, d = 0%, 2%, and 5%, and choose t_{tech} = 2005 similarly to LSH. The first value reproduces the abatement costs of \$50 and \$200 per ton of carbon avoided, which are the costs used in our previous work. The second and third values are crude representations of the potential that new technologies such as high-efficiency gas turbines, fuel cells, renewable energy sources, and 'safe' nuclear power might have on the costs of abating the emissions of greenhouse gases. For comparison, it is useful to note that the real price of electricity to U.S. consumers dropped an average of 3% per year between 1935 and 1970 due, in large part, to the effects of the innovation inherent in increasing economies of scale (Schurr et. al., 1990).

Table I:
Total-capital and annual-operating costs of emitting, low cost non-emitting, and high-cost non-emitting equipment.

j	Type of equipment	Annual operating cost σ_j (\$/kw)	Total capital cost $10\,\kappa_j$ (\$/kwh)
0	Emitting	0.025	0.152
1	Low-cost non-emitting	0.028	0.170
2	High-cost non-emitting	0.037	0.220

Figures 5a and 5b show the annual abatement costs, as a percentage of gross world product (GWP), for the ‘Do-A-Little’ and ‘Emissions-Stabilization’ abatement policies, respectively, for each of three values of the innovation parameter, relative to the ‘Conservation-Only’ case. The time-dependent gross world product is taken in this study as proportional to the $B_{CO_2}(t)$ in Eq. (1). The reader may note the thirty-year period oscillations in Figure 5b, peaking in 2090 and 2120. These are due to our assumption that all plants have an identical thirty-year lifetime. Under the ‘Emissions-Stabilization’ policy, the construction of high-cost non-emitting equipment surges around 2060, when the availability constraint in Eq. (10) first becomes important. Construction must surge again in 2090 and 2120 to replace the mass retirements of high-cost plants built thirty years before.

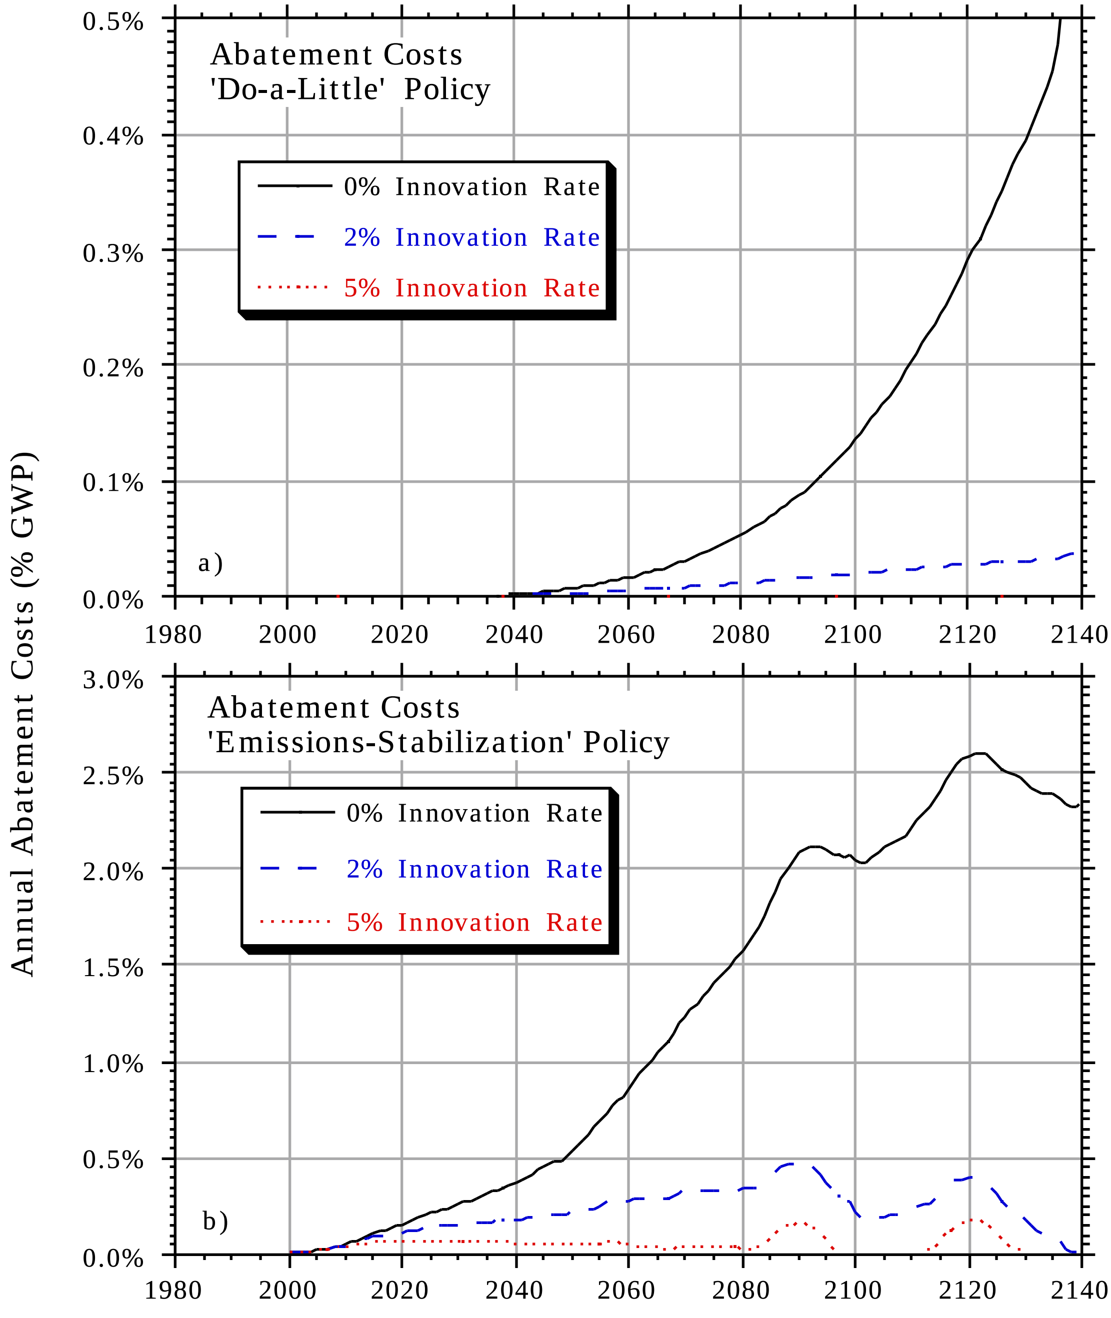

Figure 5: Annual abatement costs for the a) 'Do a Little' (∞/100/2035) and b) 'Emissions Stabilization' (30/100/2015) policies, for the technology innovation rates, d = 0%, 2%, and 5%.

3.4. Damage Model

There is little consensus among experts whether the damages due to climate change are likely to be large or small. For instance, Nordhaus (1994b) surveyed 19 scholars who had written on the economics of climate – 10 economists, four other social scientists, and five natural scientists and engineers – about the damages they expect from climate change. The respondents' 'best-guess' estimates of the damage due to a 3°C temperature rise over the next century ranged from negligible to a 21% loss in gross world product. Some respondents placed the likelihood of damages equivalent to the Great Depression (the loss of approximately one-quarter of global output), as high as 30%, while others considered the likelihood to be less than 0.5%. Such a large range of uncertainty is also found by other researchers (Dowlatabadi and Morgan, 1993), and is not surprising given our primitive understanding of the climate system, the biosphere, and society's relationship to these systems.

Given this large uncertainty, most researchers who need a quantitative expression for global-aggregate damage from climate change write the function as a quantitatively convenient power law in the global-mean surface temperature increase, $\Delta T(t)$, and choose the parameters for these functions to reproduce estimates of the aggregate global damage due to climate change (Nordhaus, 1994a, 1992; Manne and Richels, 1992; Peck and Teisberg, 1993). We follow this pattern here and express the damage in year t as a percentage of GWP as

$$D(t) = \alpha \left[\frac{\Delta T(t)}{3}\right]^3 + \beta A[\Delta T(t)] \quad . \tag{12}$$

The first term on the right is the cubic damage function used by Peck and Teisberg (1993, 1992). (Nordhaus (1994a, 1992) uses a quadratic function.) This function represents damages that rise relatively slowly in the near-term, but become relatively large towards the end of the next century. The second term on the right represents an abrupt, near-term increase in the damage. We include an abrupt-change term in the damage function because our work in LSH suggests that such phenomena could have significant effects on the long-term costs of different policies (though not on near-term policy choices). Since there is little information available to guide the choice of the potential abrupt changes in the damage function, we choose a convenient triangular function given by

$$A[\Delta T] = \begin{cases} 0 & \text{for } \Delta T < 1^{o}C \\ \Delta T - 1^{o}C & \text{for } 1^{o}C \leq \Delta T < 2^{o}C \\ 3^{o}C - \Delta T & \text{for } 2^{o}C \leq \Delta T < 3^{o}C \\ 0 & \text{for } \Delta T \geq 3^{o}C \end{cases} \quad . \qquad (13)$$

The coefficients α and β in Eq. (12) govern the size of the damage and whether the damage increases most rapidly in the near-term or long-term. We choose four values of the first coefficient, α = 0.5%, 2%, 3.5%, and 20%. The last three values correspond to the mean, median, and high estimates, respectively, from Nordhaus' survey of experts on the loss in GWP due to a 3°C rise in global-mean temperature by 2090. The first three values correspond to those used by Peck and Teisberg (1993) in their analysis of climate-change uncertainties.

For each value of α we choose two β values, β = 0 and $\alpha/2$. The first choice represents no abrupt changes, while the second produces damage functions that rapidly approach their 3°C level by the mid twenty-first century before leveling off. The resulting eight damage functions are shown for climate sensitivity ΔT_{2x} = 2.5°C and the 'Conservation-Only' abatement policy in Figure 6a, and for the 'Emissions-Stabilization' abatement policy in Figure 6b. Note that the cases with the smaller values of α and no abrupt changes provide damage functions that are similar to those used by Nordhaus (1994a, 1992) and Peck and Teisberg (1993, 1992). The abrupt-change scenarios give non-convex damage functions, which are not usually treated in optimization calculations because these functions' regions of declining marginal cost can lead to multiple locally optimum policies. There is, however, no physical barrier to non-convex damage functions. We chose to begin the abrupt change at 1°C since the mechanism for the change might reasonably be unobserved at the present level of warming, and we chose to end the effects of the abrupt change at 3°C to facilitate fitting to the estimates in Nordhaus' survey.

These damage functions represent economic and non-economic damages. While society's uncertainty about potential climate-change damages is so large that any inferences should be made with caution, the range of behaviors explored here is sufficiently large, as we will show in the next section, to call for optimum emissions abatement policies that span the full range of options from no abatement to the maximum allowed abatement. This ensemble of damage functions thus allows us to draw some important conclusions about the appropriate choice of climate-change policies.

4. COMPARING THE PERFORMANCE OF POLICY ALGORITHMS

How should policy-makers respond to the highly uncertain, but potentially serious prospects of climate change? We explicitly include the decision-maker as part of the linked system of models shown in Figure 2. Using these models, we compare the performance of three types of policy algorithms that differ in the assumptions they make about our ability to predict the future and the way they treat learning from observations.

'Optimum policies' are the best-possible policy that we would choose if we had perfect information about the future. 'Best-estimate' policies are long-term prescriptions (generally a century or more) for greenhouse-gas-emission trajectories, that minimize the expected value of long-term costs based on assumptions about the likelihood of different plausible futures. Many climate-change studies focus on either optimum or best-estimate policies, as do most debates over climate-change policy. 'Adaptive strategies' consider midcourse corrections based on observations of the climate and economic systems. As discussed above, adaptive strategies are a generalization of the sequential-strategy approach in our previous work (LSH, HLS) and the 'Learn, then act' framework of Manne and Richels (1992). Sequential strategies can make midcourse corrections at some fixed time in the future based on new information gained independently of any changes in the climate system. Adaptive strategies can gain new information based on observations, so that the rate of learning can depend on the rate of change in the climate system. We will now show that for almost any set of expectations about the future, a very simple adaptive strategy performs on average better than either of the best-estimate policies currently advocated in today's climate-change debate.

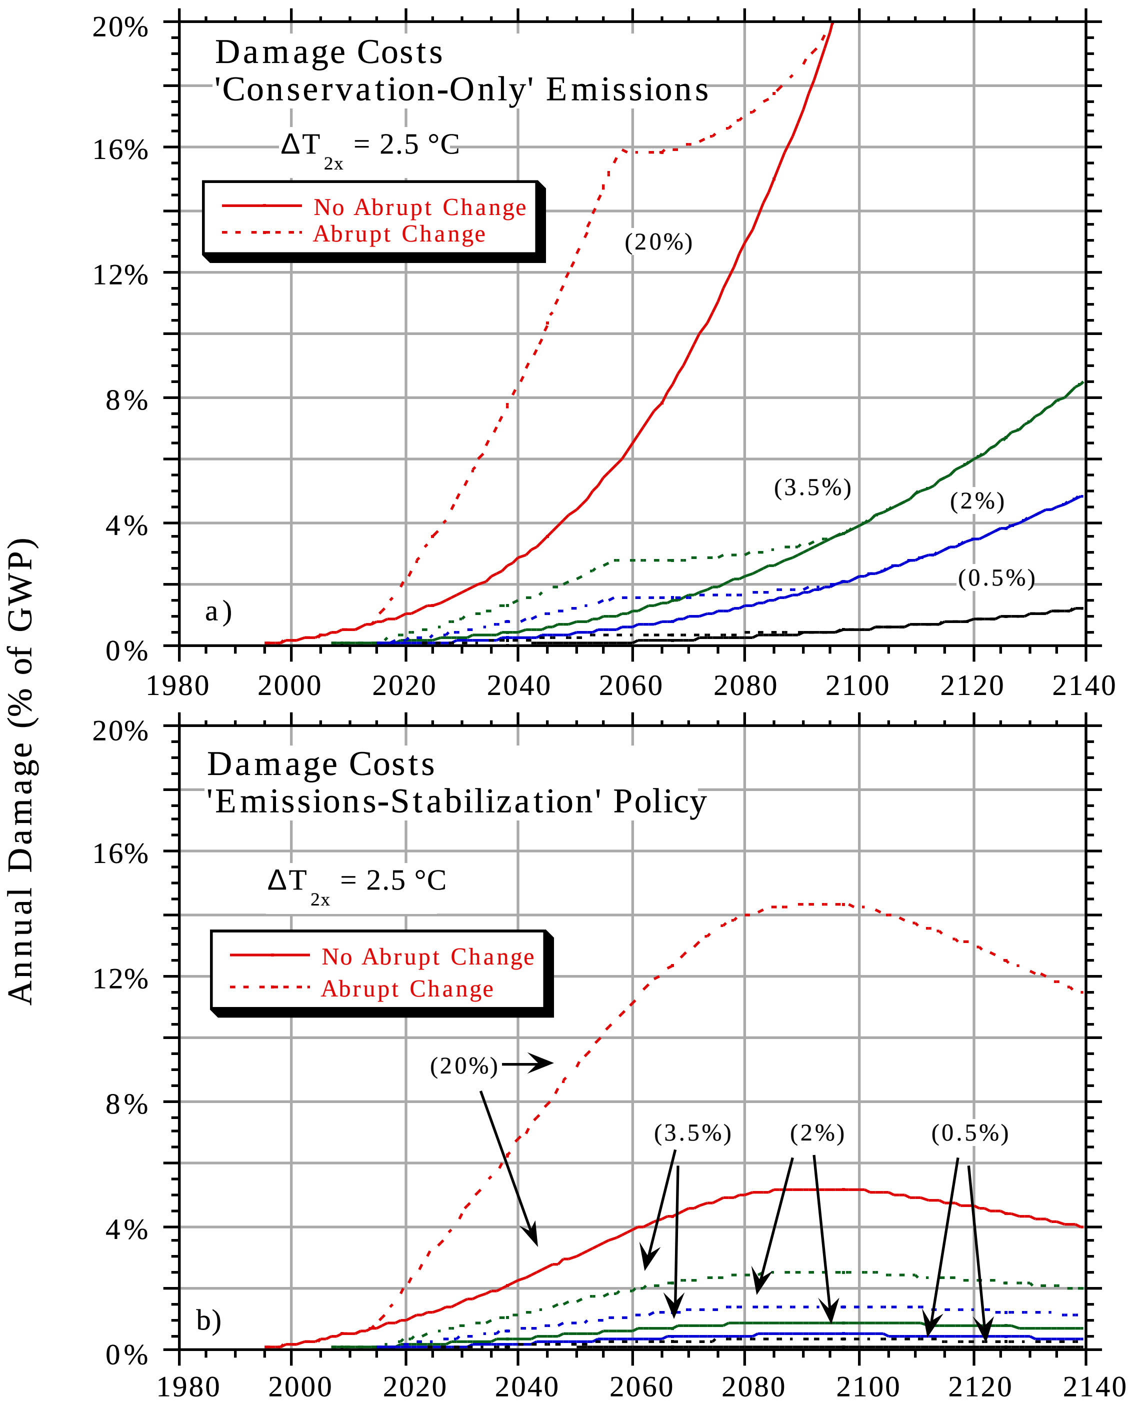

Figure 6: a) Annual damage costs for the a) 'Conservation Only' (∞/ ∞/na) and b) 'Emissions Stabilization' (30/100/2015) policies, with climate sensitivity ΔT2x = 2.5°C, and for the eight damage functions defined by a = 20%, 3.5%, 2%, and 0.5%, with (b = a/2) and without (b = 0) abrupt changes in the damage.

4.1. Optimum Policies

The uncertainties presented in Section 3 describe a three-dimensional uncertainty space in which we sample 96 points. At each point in this space there is an optimum abatement policy that produces the minimum present value of the sum of the abatement and damage-cost trajectories. Equation (4) provides abatement policies with three degrees of freedom – a first-period abatement rate, $1/R_1$; a second-period abatement rate, $1/R_2$; and the year in which we change from the first-period to the second-period rate, T_{12}. These two-period abatement policies are similar to those used in LSH and HLS. However, in our previous work we considered only two values of R_1 (40 years and infinity), and T_{12} was fixed at ten years. Here the later quantity is allowed to vary.

We examine all possible (R_1, R_2, T_{12}) triplets to find the optimum abatement policy at each of the 96 points in the uncertainty space.[8] We allow R_1 and R_2 to take any of the values 20, 30, 40, 50, 60, 75, 100, 150 years, or infinity, and T_{12} to take any of the values 2005, 2015, 2025, 2035, or 2045. Figure 7 shows the annualized global cost (abatement plus damage costs) of the optimum policy, using a discount rate of 5%, at each point in the uncertainty space. The optimum cost varies by three orders of magnitude as a function of the climate sensitivity ΔT_{2x} and the damage function (α,β). For any given combination of (α,β) and ΔT_{2x}, an increasing innovation rate, d, can reduce the optimum-response cost. The amount varies across the uncertainty space. In some cases the reductions for d = 5% approach 50 percent.

Figure 7 indicates three specific points in the uncertainty space, States of the World 1, 2, and 3, which will be a focus of our narrative. State 1 has a climate sensitivity of ΔT_{2x} = 2.5°C, damage function (α,β) = (2%,0%), and no cost reductions due to innovation (d = 0%). The optimum policy for this state is (R_1, R_2, T_{12}) = (∞/100/2035), which begins with no fuel switching in the first period and changes to a 100-year, second-period fuel-switching rate in 2035 (See Table II). The annualized global cost of this optimum policy is \$92 billion. The optimum policy for State 2, with climate sensitivity ΔT_{2x} = 2.5°C, damage function (α,β) = (3.5%,0%), and cost reductions due to innovation at a rate d = 2%, is the policy (50/100/2045), which essentially stabilizes emissions at 1990 levels until the middle of the twenty-first century (See Figure 8). This policy results in a annualized cost of \$119 billion. State 3, having ΔT_{2x} =

[8] In general exploratory modeling calls for an iterative sampling strategy to concentrate simulation runs in selected regions of the uncertainty space (Bankes, 1994 and 1993). In this work, however, the computational demands did not require us to go beyond a brute force sampling strategy.

2.5°C, (α,β) = (20%,10%), and d = 5%, represents a catastrophic climate-change threat. The optimum response for this state is the immediate and draconian emission-reductions policy (20/20/na), which results in an annualized cost of $506 billion.

Figure 8 shows the trajectory of carbon-dioxide emissions produced by the optimum policies for States 1, 2, and 3. We see that the optimum emissions trajectory for State 1 is similar to those reported by other studies that examine similar states of the world. Peck and Teisberg (1992) report an optimum carbon-dioxide emissions trajectory for a 2.5°C climate sensitivity and a cubic damage function that reaches 2% of GWP when global-mean temperature has increased 3°C. Given these similar assumptions, our optimum emissions trajectory is similar to Peck and Teisberg's through 2090. The differences before 2090 appear to be largely due to differences in assumptions about unconstrained (basecase) emissions and our assumption that emissions can be reduced 30% below the IPCC/90 Scenario A by the diffusion of low-cost conservation technologies. The differences after 2090 are likely due to Peck and Teisberg's projection that even without emission-control policies, greenhouse-gas emissions will begin to slow at the start of the twenty-second century, as society depletes its reserves of coal. Nordhaus (1994a, 1992) reports optimum carbon-dioxide emission trajectories as a fraction of uncontrolled emissions for a state of the world having 2.5°C climate sensitivity and a quadratic damage function that reaches 1.3% of GWP at a rise in global-mean temperature of 3°C. Figure 8 shows the IPCC/90 Scenario A basecase emissions used in this study, reduced by the optimum fraction reported by Nordhaus. The differences with our optimum emissions in State 1 are consistent with Nordhaus' smaller damage function.

The optimum emissions trajectories we find for States 2 and 3 do not appear to have precedents in the Integrated Assessment literature, since other studies have not predicted optimum emissions trajectories for the large damage functions and possible effects of innovation that we consider here (Parson, 1994). As expected, we find that optimum policies with relatively large emission reductions are found for states having large climate sensitivities and damage functions, which increase the potential costs of climate change, and/or significant innovation, which reduces the costs of emission reductions. For instance, for State 2, the damage function reaches 3.5% of GWP with a rise in global-mean temperature of 3°C and innovation reduces abatement costs by 2% annually, so that the optimum policy holds emissions close to 1990 levels until the middle of the twenty-first century. For State 3, the damages reach 20% of GWP with a 3°C temperature increase, so the optimum policy requires even more aggressive emission reductions.

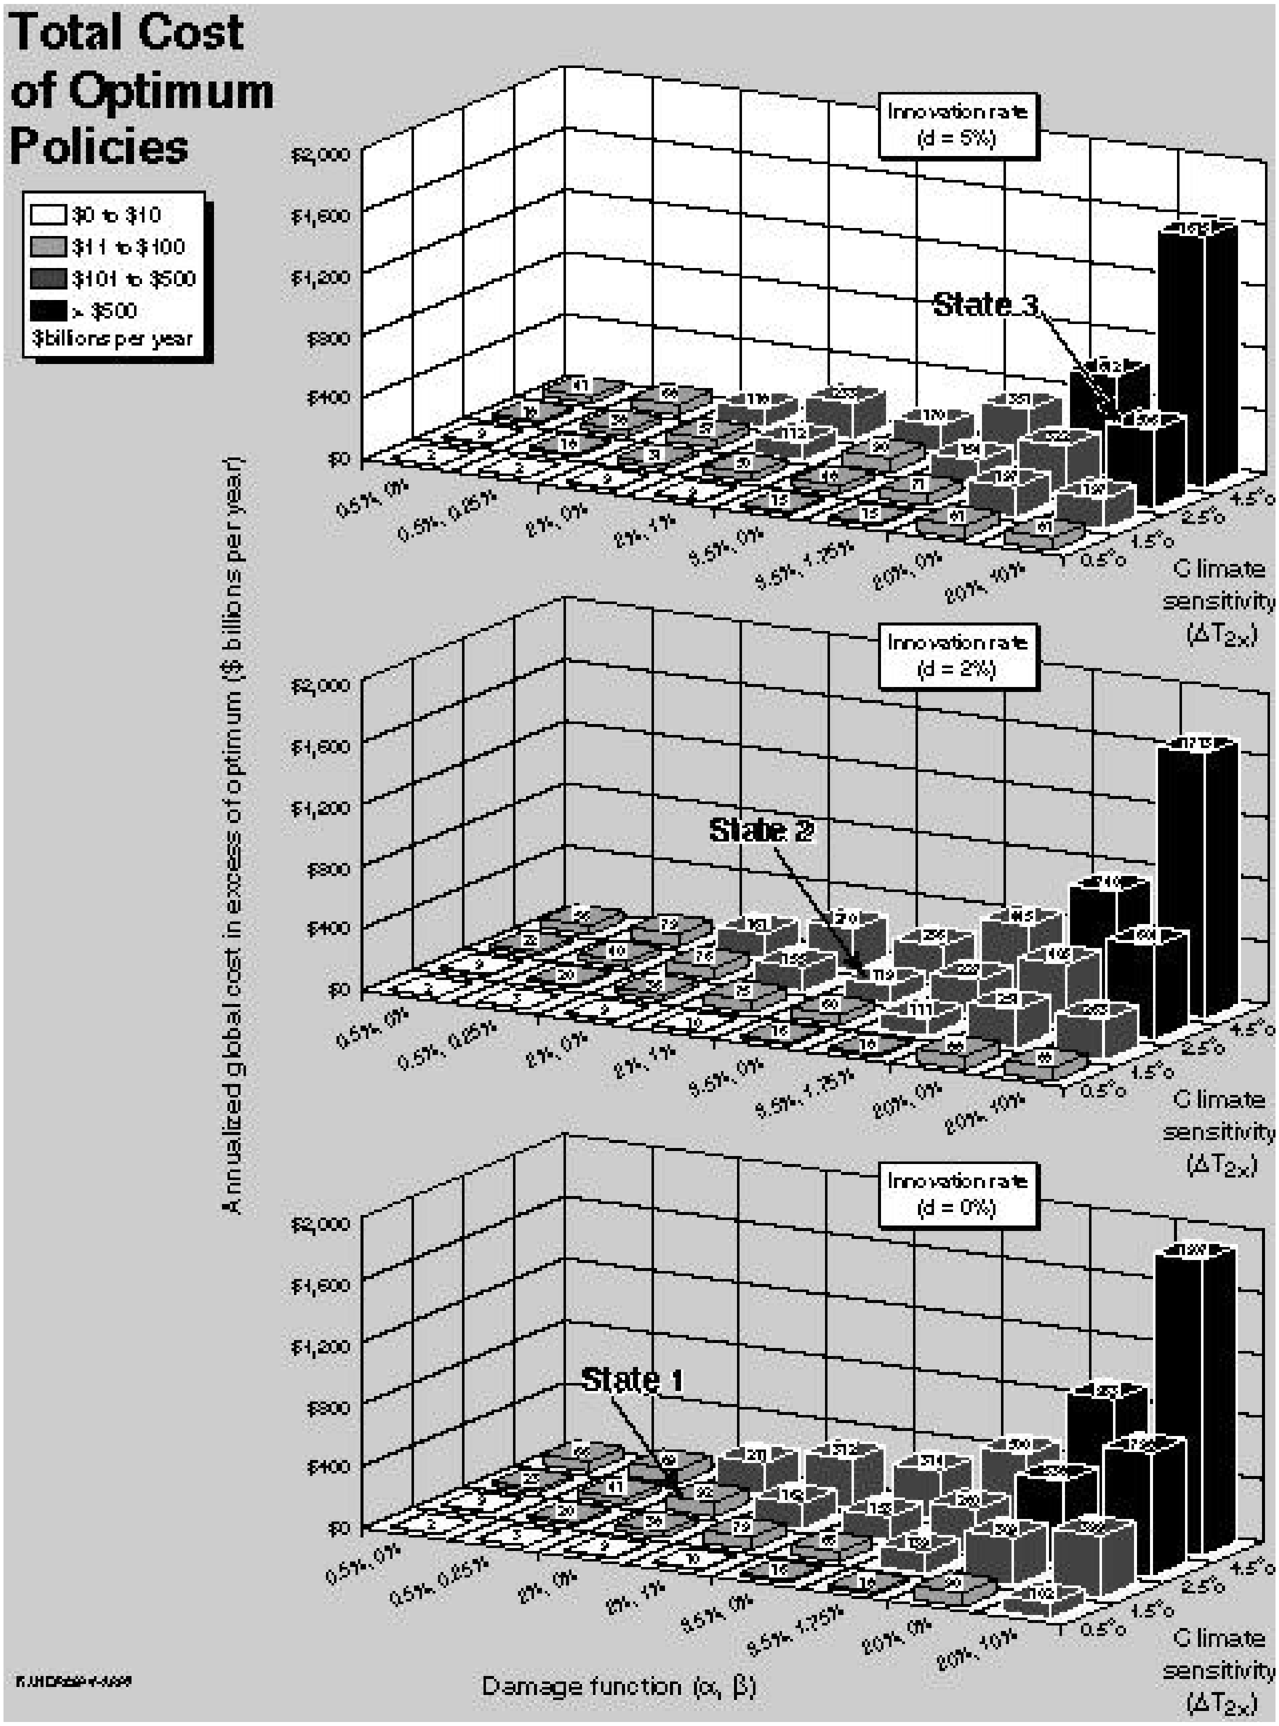

Figure 7: Annualized global cost (present value of the damage and abatement costs) of the optimum policy at each point in the uncertainty space defined by four values of the climate sensitivity ΔT2x = 0.5, 1.5, 2.5, and 4.5°C; three values of the technology innovation rate d = 0%, 2%, and 5%; and the eight damage functions (a,b) = (0.5%,0%), (0.5%, 0.25%), (2%, 0%), (2%, 1%), (3.5%, 1.75%), (3.5%, 1.75%), (20%, 0%), and (20%, 10%).

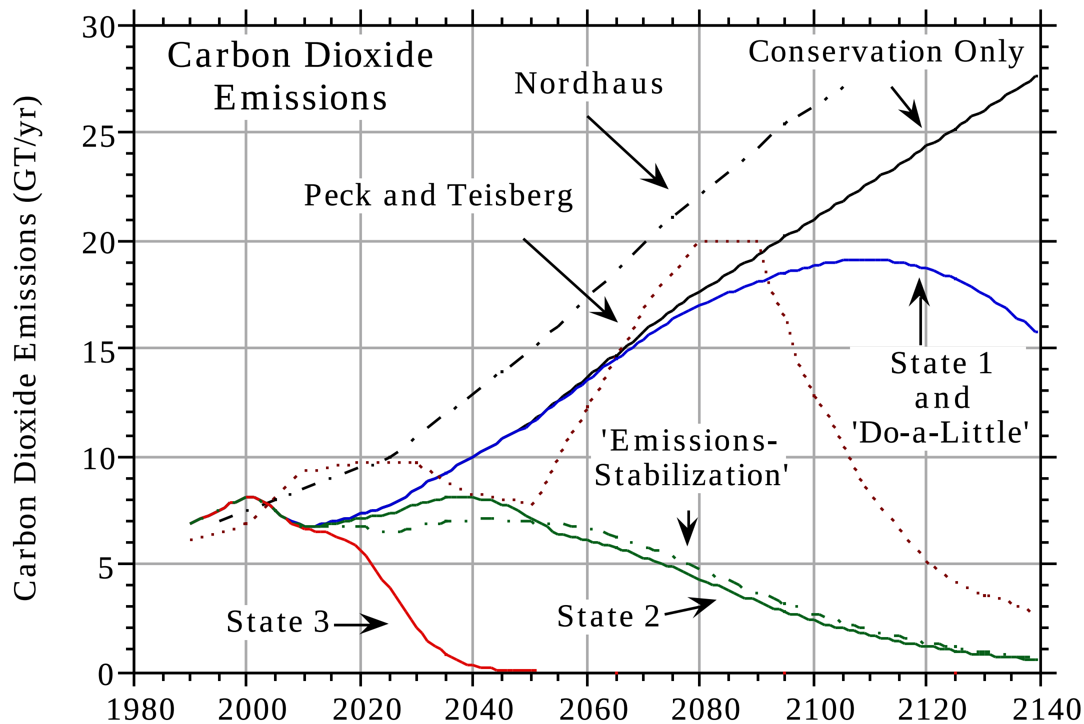

Figure 8: Optimum carbon-dioxide emissions for three points in the uncertainty space: State 1: [ΔT2x, (a,b), d] = [2.5, (2%,0%), 0%]; State 2: [2.5, (3.5%, 0%), 2%]; and State 3: [2.5, (20%, 10%), 5%], compared to the 'Conservation-Only' and 'Emissions-Stabilization' policies, and optimum emission trajectories reported by Nordhaus (1992) and Peck and Teisberg (1992) for scenarios similar to State 1.

Figure 9 shows the optimum first-period abatement rate ($1/R_1$) across the uncertainty space (Table II shows the full (R_1, R_2,T_{12}) triplet for each optimum policy). Two regions dominate the figure – one in the region of small values of the damage-function parameters (α,β) and the climate sensitivity ΔT_{2x} where the optimum policy calls for no near-term abatement ($R_1 = \infty$), and the other in the region of large damages and large sensitivity where the optimum policy calls for the fastest allowed value for near-term abatement (R_1 = 20 years).[9] States 1 and 3 exemplify these two regions. Between these two regions is a transition area, which includes State 2, where the optimum near-term emissions varies between R_1 = 150 years to R_1 = 30 years. The wide variation in optimum policy shown in Figure 9 suggests that our uncertainty space captures much of the range of disagreement between the contending camps in the climate-change debate and, conversely, that much of the divergence in opinion about climate-change policy has its roots in the wide differences in expectations about the likelihood of large

[9] There is currently little information available on the maximum rates at which society could reduce emissions. The results in Figure 9 suggest that research on this topic could be very useful in understanding the ability of society to respond to extreme climate-change scenarios.

damages[10] and/or society's ability to develop low-cost responses to the climate-change problem. (Lave and Dowlatabadi, 1993 also make this point.)

TABLE II:

Optimum policy for each point in the uncertainty space. Optimum policies for the states of the world discussed in the text -- State 1 (∞/ 100/ 2035), State 2 (50/ 100/ 2045), and State 3 (20/ 20/ na) -- are underlined. The 'Do-a-Little' and 'Emissions-Stabilization' best-estimate policies discussed in the text are (∞/ 100/ 2035) and (30/ 100/ 2015), respectively.

		Optimum Policy (R1/ R2/ T12)			
Damage	Tech	Sensitivity 0.5	Sensitivity 1.5	Sensitivity 2.5	Sensitivity 4.5
(0.5%,0%)	0%	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ ∞/ na)
(0.5%,0.25%)	0%	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ ∞/ na)
(2%,0%)	0%	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ 100/ 2035)	(50/ ∞/ 2045)
(2%,1%)	0%	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ 100/ 2045)	50/ ∞/ 2045)
(3.5%,0%)	0%	(∞/ ∞/ na)	(∞/ ∞/ na)	(50 /∞/ 2045)	(20/ ∞/ 2015)
(3.5%,1.75%)	0%	(∞/ ∞/ na)	(∞/ ∞/ na)	(20/ ∞/ 2015)	(20/ ∞/ 2015)
(20%,0%)	0%	(∞/ ∞/ na)	(40/ ∞/ 2035)	(20/150/ 2015)	(20/ 30/ 2025)
(20%,10%)	0%	(∞/ 75/ 2025)	(20/ 100/ 2015)	(20/ 20/ 2005)	(20/ 20/ 2005)
(0.5%,0%)	2%	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ 60/ 2045)	(∞/ 40/ 2025)
(0.5%,0.25%)	2%	(∞/ ∞/ na)	(∞/ ∞/ na)	(∞/ 60/ 2045)	(∞/ 40/ 2025)
(2%,0%)	2%	(∞/ ∞/ na)	(∞/ 60/ 2035)	(∞/ 50/ 2015)	(30/ 75/ 2015)
(2%,1%)	2%	(∞/ 100/ 2045)	(∞/ 60/ 2005)	(∞/ 60/ 2025)	(60/ 60/ 2005)
(3.5%,0%)	2%	(∞/ ∞/ na)	(∞/ 50/ 2025)	(50/ 100/ 2045)	(20/ 150/ 2015)
(3.5%,1.75%)	2%	(∞/ 60/ 2045)	(30/ 75/ 2015)	(20/ 100/ 2015)	(20/ 30/ 2015)
(20%,0%)	2%	(∞/ 60/ 2035)	(30/ 60/ 2015)	(20/ 60/ 2015)	(20/ 30/ 2045)
(20%,10%)	2%	(∞/ 60/ 2025)	(20/ 100/ 2015)	(20/ 20/ 2005)	(20/ 20/ 2005)
(0.5%,0%)	5%	(∞/ ∞/ na)	(∞/ 30/ 2045)	(∞/ 40/ 2025)	(150/ 20/2015)
(0.5%,0.25%)	5%	(∞/ 60/ 2045)	(∞/ 40/ 2025)	(∞/ 60/ 2025)	(∞/ 50/ 2005)
(2%,0%)	5%	(∞/ 75/ 2035)	(∞/ 50/ 2015)	(150/ 30/2015)	(20/ 100/ 2015)
(2%,1%)	5%	(∞/ 60/ 2025)	(100/ 30/ 2015)	(20/ 75/ 2015)	(20/ 30/ 2015)
(3.5%,0%)	5%	(∞/ 60/ 2035)	(150/ 30/ 2015)	(40/ 20/ 2045)	(20/ 60/ 2015)
(3.5%,1.75%)	5%	(∞/ 60/ 2025)	(20/ 100/ 2015)	(20/ 30/ 2015)	(20/ 60/ 2045)
(20%,0%)	5%	(150/ 30/ 2015)	(20/ 75/ 2015)	(20/ 30/ 2015)	(20/ 30/ 2045)
(20%,10%)	5%	(150/ 30/ 2015)	(20/ 75/ 2015)	(20/ 20/ 2005)	(20/ 20/ 2005)

[10] Different expectations about damages may be due to differences in the values assigned to the effects of climate change as well as to differences in predictions of the actual physical effects.

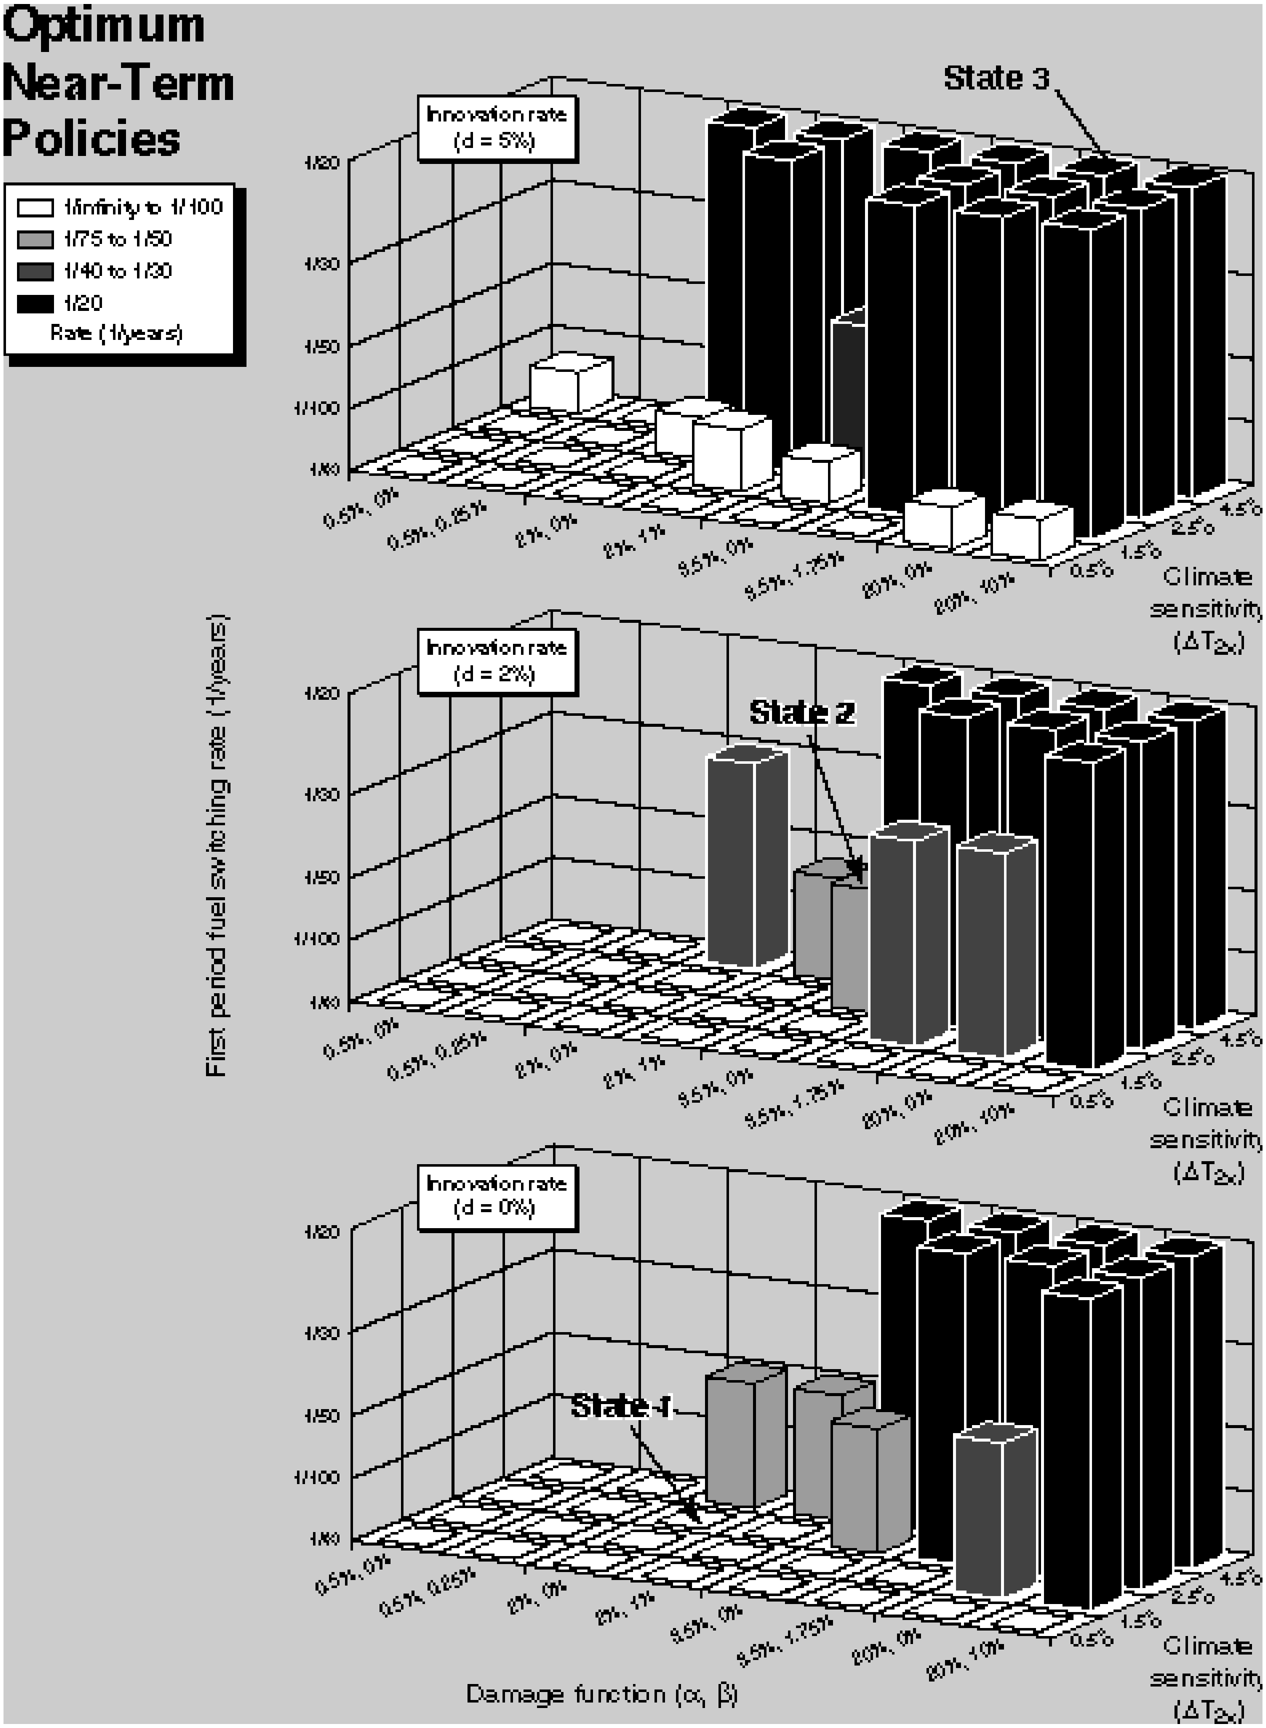

Figure 9: The (R1, R2, T12) triplets for the optimum policy at each point in the uncertainty space. Infinite values of R1 and R2 shown as 210 years.

In and around the transition area in Figure 9, the optimal near-term policy displays some interesting behavior. In some places, for instance, the optimal near-term policy is sensitive to the shape of the damage function. The state [$\Delta T2x = 2.5°C$, (a,b) = (2%,0%), d = 5%] calls for optimum near-term abatement of R1 = 150 years, while the state [$\Delta T2x = 2.5°C$, (a,b) = (2%,1%), d = 5%] calls for optimum near-term abatement of R1 = 20 years. This suggests that studies which consider only convex damage functions like those shown in Figure 1 may need to qualify statements they make about the desirability of near-term abatement.

In the transition region, the optimal policy can also shift abruptly for small changes in the state of the world. For instance, the state [$\Delta T2x = 1.5°C$, (a,b) = (3.5%,1.75%), d = 0%] calls for the optimum response (∞/∞/na), with no first-period fuel-switching, while the neighboring state [2.5°C, (3%,1.75%), 0%] calls for the opposite, an optimum policy (20/∞/2015) that begins with twenty years of the maximum fuel-switching rate. We observe this large shift because these two very different policies have similar costs in this region of uncertainty space. The policy (∞/∞/na) has no abatement costs but suffers large damage costs. In contrast, (20/∞/2015) limits damages but suffers high abatement costs. When $\Delta T2x = 1.5°C$, the abatement costs for (20/∞/2015) are sufficiently high that (∞/∞/na) is the better policy. When $\Delta T2x = 2.5°C$, the damage costs for (∞/∞/na) rise enough so that (20/∞/2015) becomes the better policy. This behavior is characteristic of non-convex and non-linear systems. For some small changes in state of the world, the optimal policy can jump a large distance in policy space, reflecting a shift in the global optimum among the multiple local peaks in the policy landscape (Kauffman, 1989). The large and not always smooth variation across the uncertainty space shown in Figure 9 suggests that, in practice, the optimum policy may be very difficult to predict with confidence.

4.2 Best-Estimate Policies

Because the uncertainties about climate change make it difficult to know the optimum policy, analysts often recommend abatement policies based on the best available estimates. Often this entails choosing the policy option with the smallest expected value of projected costs, averaged over the uncertainty space, given some estimate of the relative likelihood of each point in the space. In many circumstances, this method can provide an excellent basis for informing policy decisions. However, when the uncertainties about the future are very large, and when we possess very limited information as to the probability distribution over the uncertainty space, such best-estimate methods provide policy prescriptions that largely reflect the analysts' implicit or explicit judgments about the likelihood of events about which we have little ability to predict.

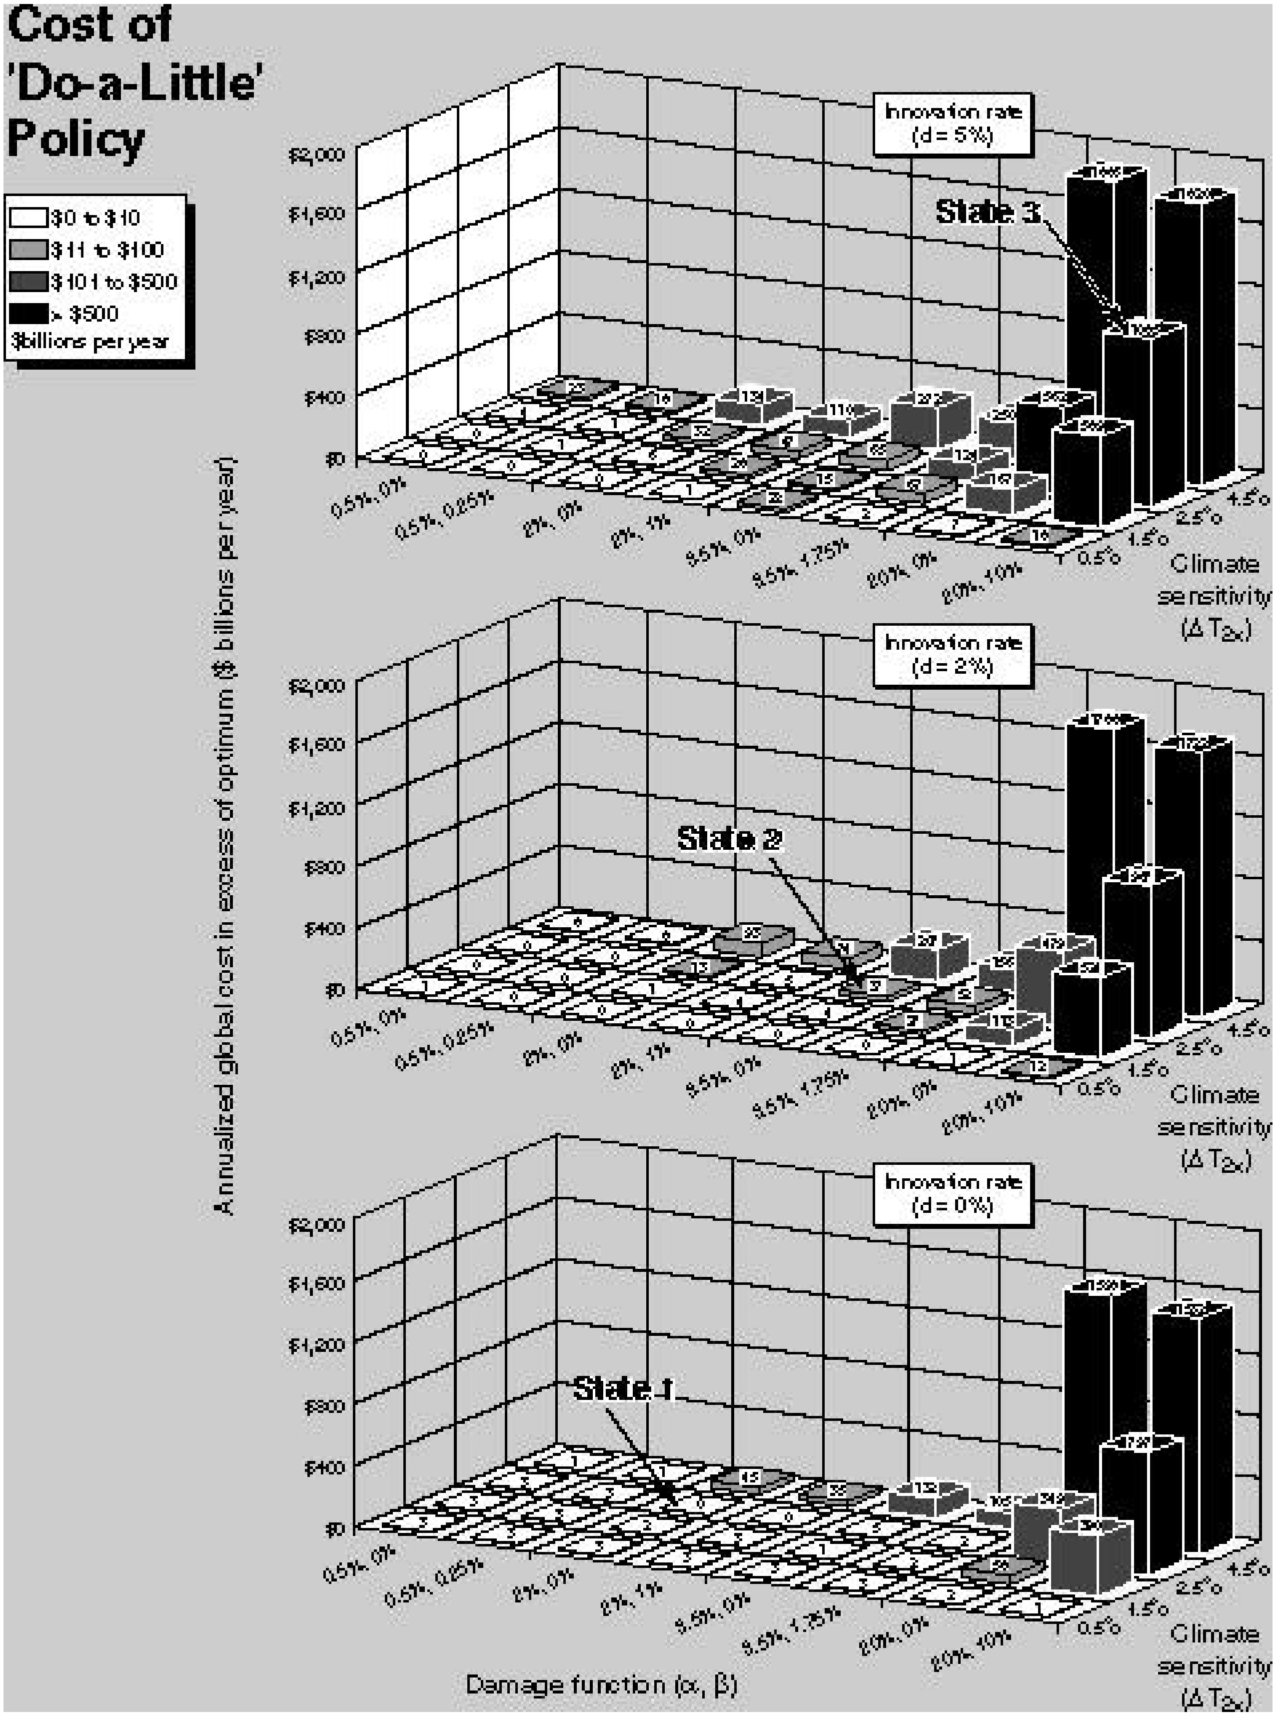

Figure 10: Annualized global cost in excess of the optimum cost of the 'Do-a-Little' policy, (∞/100/2035), at each point in the uncertainty space.

To illustrate, we examine the performance of two abatement policies similar to those that have emerged as the preferred best-estimate prescriptions among two different communities of analysts debating the climate-change problem. The ‘Do-a-Little’ abatement policy, defined here by the triplet (∞/100/2035) and shown in Figure 8, is the optimum response for State 1 and, as we have seen, is similar to the optimum policies found in many economic cost/benefit analyses of potential climate-change policies. The ‘Emissions-Stabilization' abatement policy, defined as (30/100/2015) and shown in Figure 8, is close to the optimum policy for State 2 and returns and holds global emissions to their 1990 levels through 2060. This policy is similar to that advocated by many scientists and required by many nations’ commitments under the Rio Treaty.

Figure 10 shows the difference between the annualized global costs of the 'Do-a-Little' and optimum policies at each point in the uncertainty space. Not surprisingly, the cost difference is zero or negligible in those regions of the space with small damages and climate sensitivity. However, the ‘Do-a-Little’ abatement policy makes substantial errors in those regions of the uncertainty space that are not consistent with the policy's underlying assumptions. In those regions of uncertainty space characterized by large climate sensitivities, large damages, or significant innovation, the ‘Do-a-Little’ policy can have large annualized costs up to 1.9 trillion dollars per year more than the optimum.

The ‘Emissions-Stabilization’ abatement policy, as shown in Figure 11, performs best in a region with innovation parameter d = 2% or 5% and running from low damage and high sensitivity to high damage and low sensitivity. In regions of low sensitivity and low damage, 'Emissions-Stabilization' makes errors because it is too aggressive. In regions of high sensitivity and high damage, it makes errors because it is not aggressive enough.

The choice between an irrevocable ‘Do-a-Little’ policy and an irrevocable ‘Emissions-Stabilization’ policy depends on society’s long-term estimates of the future. ‘Do-a-Little’ will hold down costs if damages are small. However, it risks huge costs if damages are large. ‘Emissions-Stabilization’ provides some insurance against large damages, but is relatively costly if damages are small or if innovation fails to bring down the cost of abatement. Figure 12 compares the expected value of the annualized costs of the ‘Do-a-Little’ and ‘Emissions-Stabilization’ policies as a function of the probabilities of extreme damages, p; significant innovation, q; and extreme climate sensitivity, s. The probabilities are defined in Table III for the parameters that define our uncertainty space. Surface A in Figure 12 separates the region where the

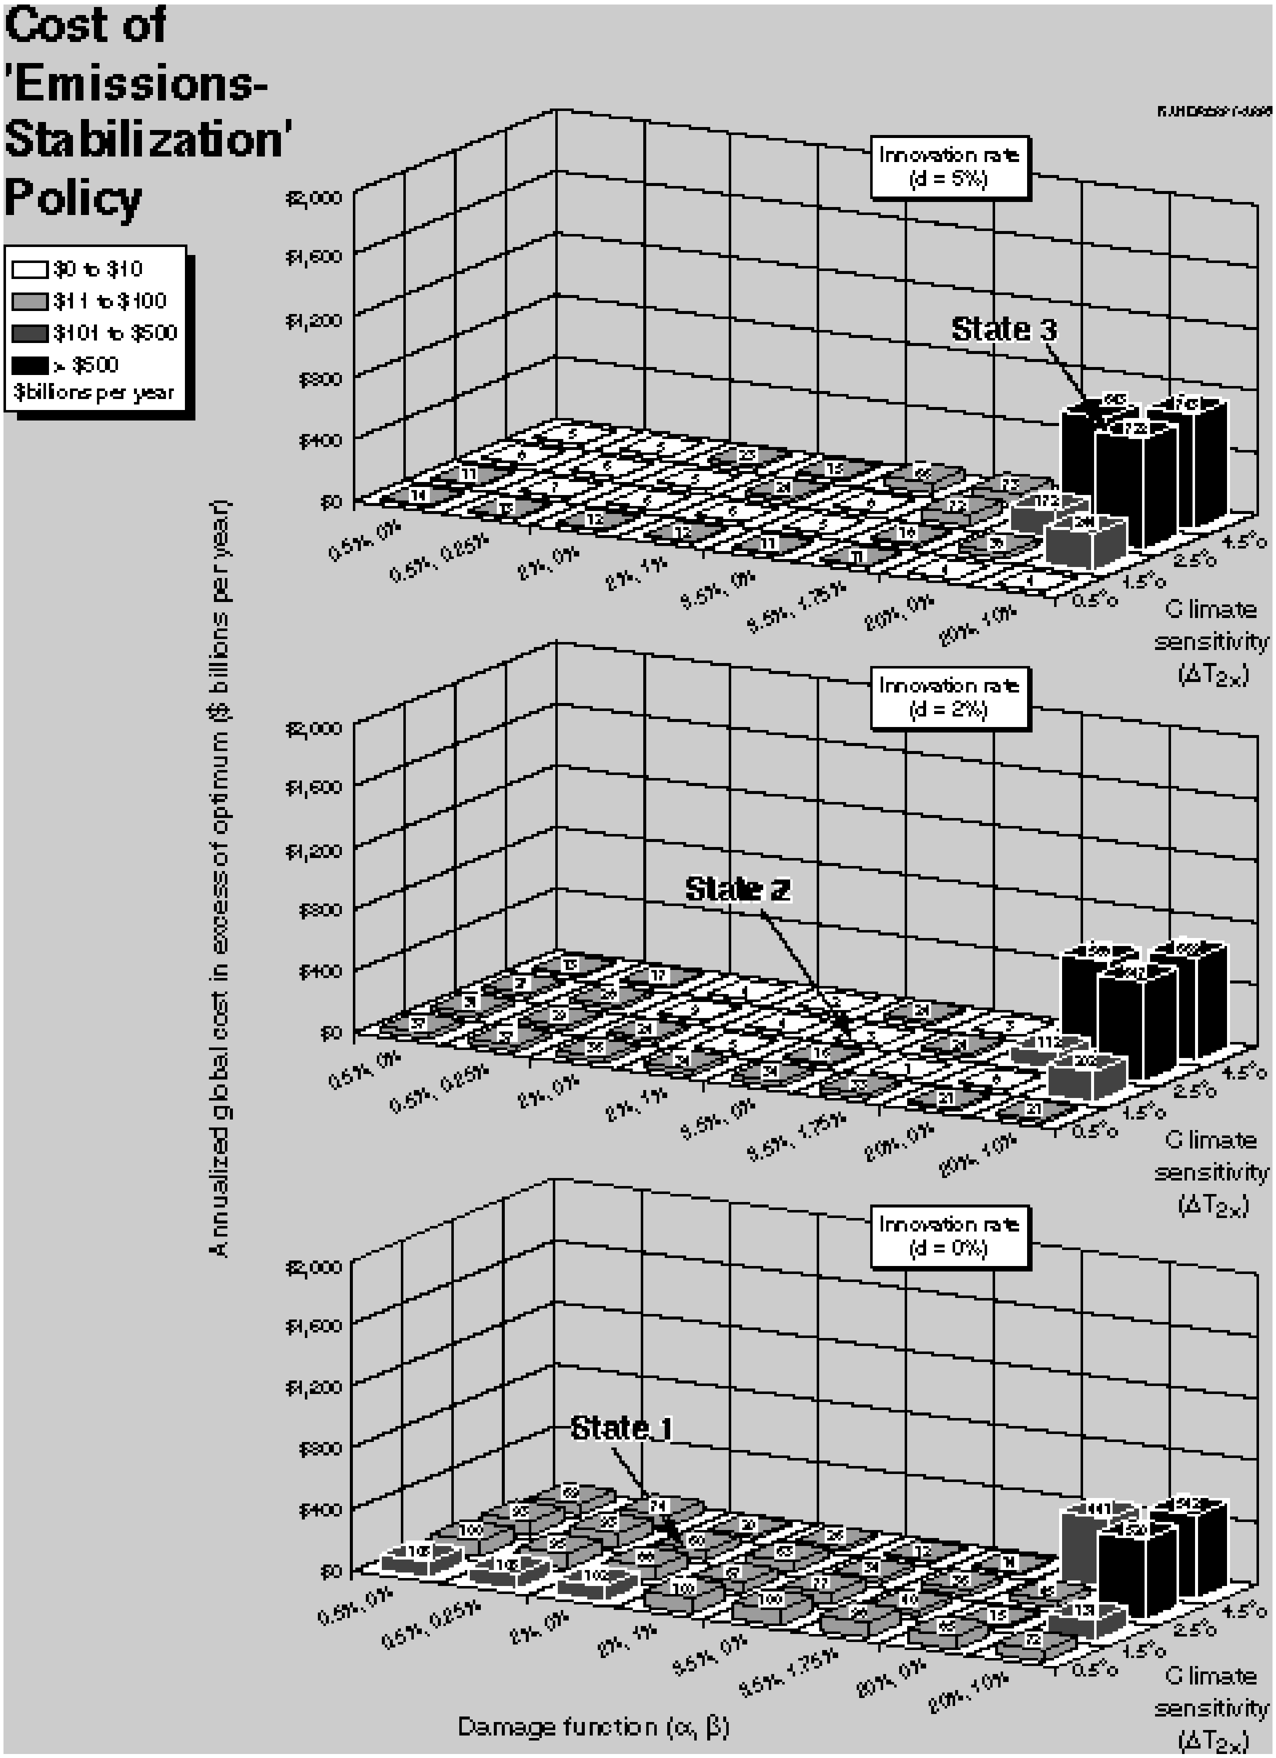

Figure 11: Annualized global cost in excess of the optimum cost of the 'Emissions Stabilization' policy, (30/100/2015), at each point in the uncertainty space.

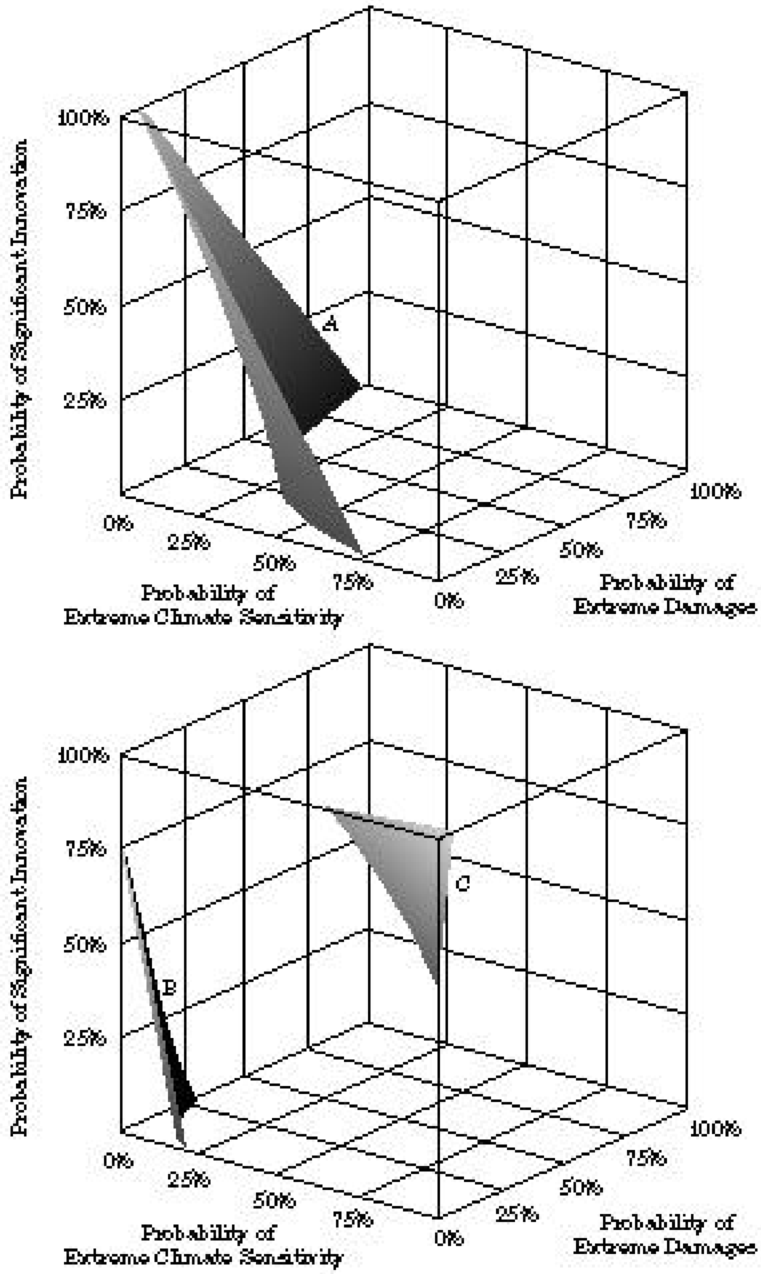

Surface A:
Do-a-Little policy preferred to the left | Emissions-Stabilization policy preferred to the right

Surface B:
Do-a-Little policy preferred to the left | Adaptive Strategy preferred to the right

Surface C:
Adaptive Strategy preferred below | Emissions-Stabilization policy preferred above

Figure 12: Surfaces separating the regions in probability space where the expected value of (A) the 'Do-a-Little' policy is preferred over the 'Emissions-Stabilization' policy, (B) the adaptive strategy is preferred over the 'Do-a-Little' policy, and (C) the adaptive strategy is preferred over 'the Emissions-Stabilization' policy , as a function of the probability of extreme damages, significant innovation, and extreme climate sensitivity, as defined in Table III.

'Do-a-Little' policy is preferred (to the left of Surface A) from the region where 'Emissions-Stabilization' is preferred (to the right of Surface A). We see that the choice between the two policies is strongly dependent on society's expectations about all three dimensions of the uncertainties considered in this study. 'Do-a-Little' is preferred over 'Emissions-Stabilization,' even for large probability of extreme damages (p ~ 95%), if society assigns zero probability to an extreme climate sensitivity and to significant innovation. (Note that in Figure 12, 'extreme damages' includes damage functions with abrupt changes, $\beta = \alpha/2$, as well as very large damages, α = 20%.) However, if society assigns even moderate probability to extreme climate sensitivity and significant innovation (s and q ~ 25%), then 'Emissions-Stabilization' can be preferred, even with moderate probability of extreme damages. If society assigns near-zero probability to extreme damages, 'Do-a-Little' is preferred unless there is high probability of extreme sensitivity or significant innovation (s ~ 75% or q > 100–4s/3%).

TABLE III:
Probabilities used in Figure 12.

	Probability of Extreme Damages:	**p**
(a,b) =	(0.5%,0%); (2%, 0%); (3,5%, 0%)	(1 - p)/3
(a,b) =	(0.5%,0.25%); (2%,1%); (3.5%,1.75%); (20%,0%); (20%,10%)	p/5
	Probability of Significant Innovation:	**q**
d =	0%	(1 - q)
d =	2%, 5%	q/2
	Probability of Extreme Climate Sensitivity:	**s**
ΔT2x =	0.5°C, 1.5°C, 2.5°C	(1 - s)/3
ΔT2x =	4.5°C	s

The particular curves shown in Figure 12, are sensitive to our choice of the extreme cases in our uncertainty space and to our choice of probability distributions in Table III. Nevertheless, the central message from Figures 10, 11 and 12 is that the choice between 'Do-a-Little' and 'Emissions-Stabilization' policies is neither a very realistic nor very useful way to frame the climate-change policy problem. This result is insensitive to the particular assumptions that underlie these figures. Both policies are the wrong choice for large portions of the uncertainty space and both impose significant costs if they are chosen incorrectly. Given the uncertainty about the future, there is a high risk of error choosing either course. Thus, if we confine ourselves to comparing best-estimate approaches, there is no way to resolve the debate between the advocates of 'Do-a-Little' and 'Emissions-

Stabilization.' This dilemma is not unique to the particular pair of best-estimate policies considered here. As is clear from an examination of the optimum policies in Table II and in Figures 7 and 9, there is no single best-estimate policy that can be successful everywhere in the uncertainty space. If forced to choose any best-estimate policy and stick to it for all time, we risk estimating incorrectly with very expensive consequences.

4.3. Adaptive Strategies

Society, of course, is not likely to choose one policy and stick to it forever regardless of subsequent events. Society will do the best it can to adapt its climate-change policy to changing circumstances. In this section we will explicitly model such an adaptive process. We will compare the performance of an adaptive strategy – one that can make midcourse corrections based on endogenous learning from observations of the climate and economic systems – to the performance of the 'Do-a-Little' and 'Emissions-Stabilization' policies considered above. We will show that the adaptive strategy can avoid the large errors made by the two best-estimate approaches and thus, on average, perform better over the uncertainty space. Given the large uncertainty associated with climate change, the adaptive strategy's success is not surprising and, in fact, many policy discussions based on best-estimate studies do focus on the implications for near-term policy choices. Thus, the purposes of the comparisons in this study are to: (i) demonstrate quantitatively that a simple adaptive strategy performs better on average than the two best-estimate policies for virtually the entire set of expectations about the future, and (ii) suggest that such a model can yield important and unique insights about the actual decisions society will be forced to make as it responds to the threat of climate change.

We examine a very simple adaptive strategy, designed explicitly to facilitate comparison with the optimum and best-estimate policies considered above. As shown in Figure 13, the strategy begins with a predetermined abatement rate, R_1 = 100 years. Once every decade starting in 2005, the strategy observes the behavior of two time series in the modeled system – the abatement costs K(t) and the damage costs D(t). If the damage (as a percentage of gross world product, GWP) exceeds a predetermined threshold,

$$D(t) > D_{thres} \tag{14a}$$

at the first observation, the strategy switches to a second-period rate R_2 = 20 years, to give the policy (100/ 20/ 2005). If either the damage costs satisfy condition (14a) or the annual change in abatement costs (as a percentage of gross world product, GWP) satisfies

$$\dot{K}(t) \equiv K(t) - K(t-1) < K_{thres} \qquad (14b)$$

at a subsequent observation, the strategy switches to a second-period rate R_2 = 40 years, to give the policy (100/ 40/ T_{12} = t). If neither condition (14a) nor (14b) is met by 2045, the strategy maintains the first-period rate indefinitely. The resulting policy is (100/100/ na).

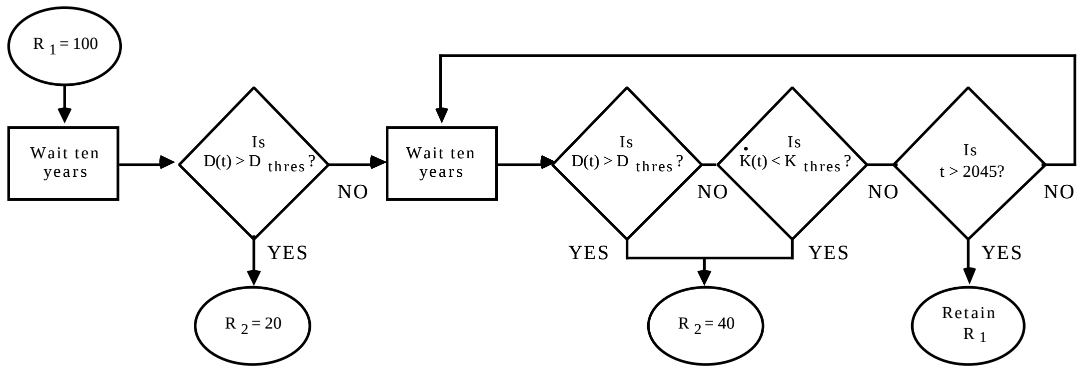

Figure 13: Flow chart describing the adaptive strategy.

It would certainly be possible to design a more sophisticated strategy that would perform better across the uncertainty space than the one used here. However, as we will show below, this simple strategy is sufficient to demonstrate that strategies that can make midcourse corrections based on observations can substantially outperform policies that do not. The first condition (14a) allows the strategy to detect and respond to rapidly rising damage. The second condition (14b) allows the strategy to detect and respond to innovation that drives down abatement costs. We selected the above-mentioned values for R_1, R_2, and set D_{thres} = 0.3% of GWP and K_{thres} = 0.0004% of GWP, after examining a variety of values and choosing those which seemed to produce, on average, the lowest costs everywhere across the uncertainty space. These threshold values are unlikely to produce the best-possible performance for the adaptive strategy, and a superior set could likely be found if we invested the computational resources to discover it.

Figure 14 demonstrates how the adaptive strategy responds to the three specific states considered in Figure 8. For the first state, the optimum response was 'Do-a-Little.' The adaptive strategy for this state starts with long-term abatement at a fuel switching rate of 100 years. Because none of the thresholds in Eq. (14) is triggered in this case, the strategy never switches to a second-period abatement rate. The strategy's emissions reductions are greater than the optimum because the conditions in Eq. (14) give it no means to check whether its first-period reductions are too aggressive. The second state, calls for an optimum emissions policy of (50/100/2045), close to 'Emissions-Stabilization.' The adaptive strategy responds in this case with a first-period abatement rate of 100 years and switches to a second-period rate of 40 years in 2035 when the rising damage triggers threshold (14a). The optimum response to the third state is immediate and draconian emissions reductions. The adaptive strategy begins with ten years of abatement at a 100-year rate, then switches to the 20-year rate when the damage threshold is triggered in 2005.

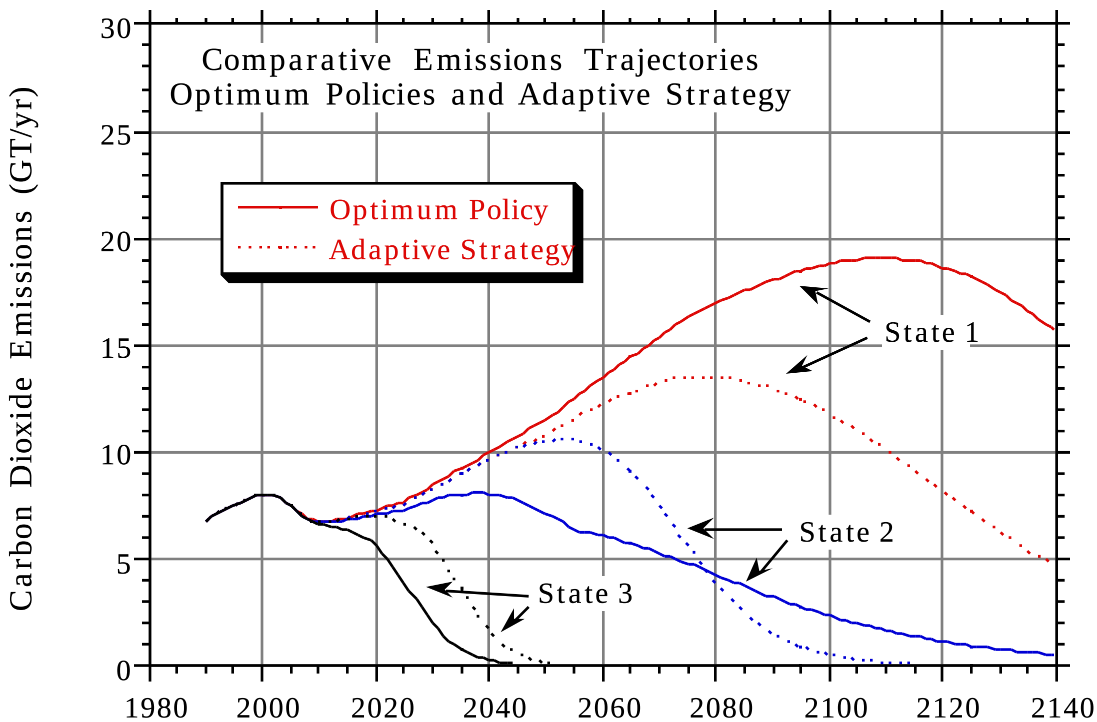

Figure 14: Time evolution of carbon-dioxide emissions for the optimum policy and the adaptive strategy at three points in the uncertainty space: State 1: [ΔT2x, (a,b), d] = [2.5, (2%,0%), 0%]; State 2: [2.5°C, (3.5%, 0%), 2%]; and State 3: [2.5°C, (20%, 10%), 5%].

In much of the uncertainty space, the adaptive strategy costs more than the best-estimate policy that works best in that region. However, the adaptive strategy never makes errors as large as the two best-estimate policies at their worst. Figure 15 shows the difference in annualized costs between the adaptive strategy and the optimum policy at each point in the uncertainty space. For the first state, the 'Do-a-Little' policy is the optimum policy, and the annualized global cost of following the adaptive strategy imposes $6 billion per year in excess costs. However, this is considerably less than the $60 billion per year excess cost of following the 'Emissions Stabilization' policy (see Figure 11) for State 1. For State 2, both the 'Emissions-Stabilization' policy and the adaptive strategy are close to the optimum policy, and cost only $3 billion per year and $5 billion per year, respectively, more than the minimum-possible cost. The 'Do-a-Little' policy, however, is not a successful response to State 2 and imposes $37 billion per year in costs over the optimum policy (See Figure 10). For the catastrophic damages of State 3, the adaptive strategy costs $265 billion per year more than the cost of the optimum policy, which is much smaller than the excess costs of $1085 and $722 billion per year for the `Do-a-Little' and 'Emissions Stabilization' policies, respectively.

Figure 12 shows the preferred choice among the adaptive strategy and the 'Do-a-Little' and the 'Emissions-Stabilization' policies, as a function of the probabilities of extreme damages, significant innovation, and extreme climate sensitivity. In contrast to the choice between 'Do-a-Little' and 'Emissions-Stabilization,' we see that the expected value of the adaptive strategy dominates that of the two best-estimate policies for almost any expectation about the uncertainties. Unless society is highly certain that over the coming decades there will be no significant innovation, no extreme damages, and no extreme sensitivity (to the left of Surface B, in the region where the product $p\ q\ s < 1/1,000$), it should prefer the adaptive strategy over 'Do-a-Little'. Society should also prefer the adaptive strategy over 'Emissions-Stabilization' unless it is almost certain that the probability of extreme damages is negligible and the probabilities of extreme climate sensitivities and significant innovation are very high (above Surface C, in the region where $p \sim 0\%$ and q and $s \sim 100\%$). 'Emissions-Stabilization' is preferred, perhaps unexpectedly, in this region of probability space because it is the only region that: (i) encompasses the band of uncertainty space where 'Emissions-Stabilization' performs better than the adaptive strategy and (ii) has negligible chance of very large damages ($\alpha = 20\%$) for which the adaptive strategy is superior to 'Emissions-Stabilization'.

Figure 12 demonstrates that the adaptive approach is superior to either of the best-estimate policies, unless the uncertainties about climate change are very small. In addition, even in the regions where the 'Do-a-Little' or 'Emissions-Stabilization' policy is preferred, the costs of erroneously choosing the adaptive strategy are relatively low.[11] The adaptive strategy helps resolve the debate between the advocates of the 'Do-a-Little' and 'Emissions-Stabilization' policies because the adaptive strategy performs better than both policies, largely independently of our expectations about the future.

[11] For any given set of expectations about the future, we could choose an adaptive policy which always performed at least as well as the best-estimate policy. The point here is that we can chose an adaptive policy which performs adequately independent of our expectations.

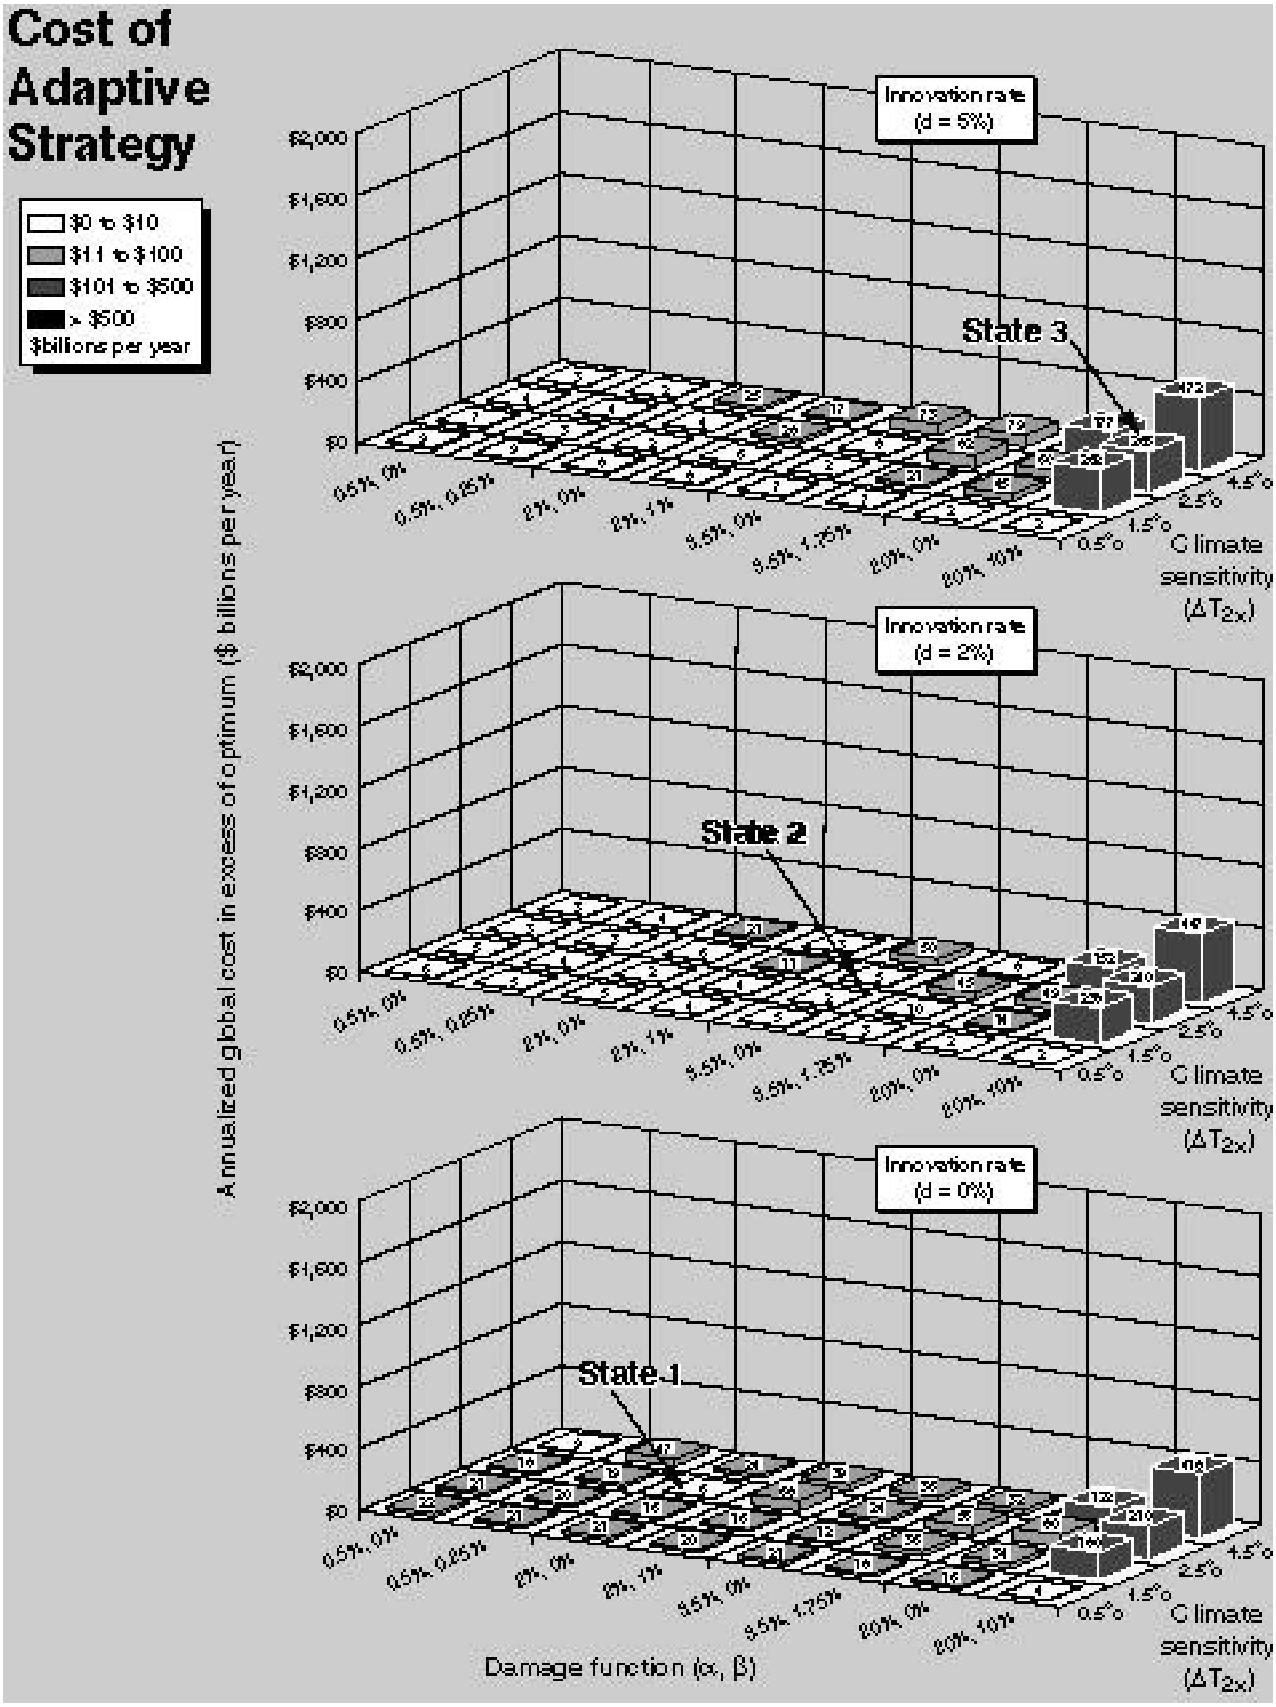

Figure 15: Annualized global cost in excess of the optimum cost of the adaptive strategy at each point in the uncertainty space.

5. IMPLICATIONS

The exploratory-modeling/adaptive-strategy approach presented in this study has allowed us to compare the performance of a simple adaptive strategy and the 'Do-a-Little' and 'Emissions-Stabilization' best-estimate policies across a vast uncertainty space. We find that the choice between the two best-estimate policies is strongly dependent on society's expectations about all three uncertainties considered in this study – the probabilities of extreme damages, extreme climate sensitivities, and significant innovation. Much of the divergence in opinion between advocates of the 'Do-a-Little' and 'Emissions-Stabilization' policies may have roots in widely different expectations about the likelihood of these factors. If so, best-estimate approaches may have difficulty producing robust policy recommendations as long as society lacks the ability to make accurate, long-term predictions of the climate and economy.

We also find that even the simple adaptive strategy considered here dominates the performance of the best-estimate policies for virtually the entire set of expectations about the future. This result suggests that, even without accurate predictions, policy-makers should be able to craft an adaptive strategy towards climate change in which society explicitly plans to make midcourse corrections based on observations of the climate and economic systems.[12] In addition, the methodology presented here may help us to better frame and address the questions that will arise if policy-makers pursue such an adaptive strategy.

One important question that emerges from the adaptive-policy framework is whether noise, climate oscillations, measurement errors, and our imperfect understanding of the climate and economic systems will allow an adaptive strategy to extract meaningful information from the phenomena it observes. This question is central to the performance of an adaptive strategy, but does not arise at all in the traditional best-estimate approach. For instance, the strategy used in this study responds to observations corresponding to annual global damages due to climate change of about $100 billion and annual global changes in abatement costs of $100 million per year (Eq. 14). It is not at all clear whether such unambiguous measurements would be simple to make, or impossible. As an indication of the scale of these numbers, note that the 1982-83 El Niño event caused worldwide damages estimated to be about $8 billion (NCAR, 1994). The more ambiguous its measurements, the more an adaptive strategy will be prone to errors of

[12] This result need not be true for all problems with large uncertainty. Rather this result derives from our assumptions about the relative rates of change of the climate and economic systems. In some problems, it may be impossible to make timely responses to observations so that society can do no better than pursuing a best-estimate policy.

commission and omission, adopting aggressive abatement when it is not warranted and not adopting aggressive abatement when it is warranted. If its observations are sufficiently poor, an adaptive strategy will eventually perform no better than a best-estimate policy. The methodology proposed here appears to provide a suitable framework for addressing these key questions about the ambiguity of observations. For example, the adaptive-strategy approach would quickly focus research attention on the question of how we might detect the onset of extreme climate damages.

The adaptive framework may also help us recast the climate-change debate from the current argument over the appropriate levels of near-term abatement to encompass a wider set of policy choices. In this light, it is worthwhile to note some important similarities between the analytic framework for adaptive strategies and some of the strategic-planning techniques increasingly being employed by many successful business firms and other organizations in recent years as they attempt to cope with very uncertain and turbulent times. For instance, the scenario-planning techniques developed by Royal Dutch Shell (De Geus, 1988; Wack, 1985a, 1985b) and the Assumption Based Planning developed for the U.S. Army (Dewar et al., 1993), ask a much broader array of questions than generally addressed in the current climate-change debate: How might the assumptions that undergird current policies fail? What observations will suggest the need for a change of policy? How can we hedge to improve our array of available options if the need to change current policy arises?

Our adaptive-strategy approach provides an analytic framework to address many of these issues. We identify regions of uncertainty space where different policies fail, and examine the observations that should cause a change in policy. While we do not consider hedging strategies in this work, they could be included in our approach as policies that, for example, improved the options available to the adaptive strategy. For instance, a hedging policy might focus on research and demonstration programs for non-emitting energy technologies that would increase the likelihood that abatement costs would drop in the future.

In contrast, traditional best-estimate approaches often view hedging against high-consequence, (potentially) low-probability events as averaging the different abatement policies appropriate for the high and low estimates of climate-change damages. While these average policies improve the expected value, under many conditions they will still produce costly errors unless the actual damages turn out to be the average of the high and low estimates. An adaptive climate-change policy would view hedging as building contingency plans and responding to opportunities and dangers as they become apparent.

The adaptive-strategies framework may help policy-makers make reasonable and defensible choices about near-term climate-change policy without the need for accurate and widely-accepted predictions of the future. This approach raises new questions such as how near-term policy might encourage the development of better options for massive future emission reductions than those currently available and what observations ought to trigger such large reductions. We believe that the exploratory-modeling/adaptive-strategy approach presents a first step towards an analytic framework for addressing such questions.

Acknowledgments

The authors thank Jim Gillogly for performing the computational experiments for this study, Joe Nuzzolo for his valuable assistance with the computer code, and Ayman Ghanem for the 3-dimensional graphics of Figure 12. This research was supported in part by the U.S. National Science Foundation and the Carbon Dioxide Research Program, Environmental Sciences Division of the U.S. Department of Energy under Grant ATM-9001310, and by the RAND Corporation.

REFERENCES

Bankes, S. C.: 1993, "Exploratory Modeling for Policy Analysis", Operations Research, vol. 41, no. 3, 435-449, (also published as RAND RP-211).

Bankes, S. C.: 1994, "Computational Experiments and Exploratory Modeling" Chance, vol. 7, no. 1, 50-57, (also published as RAND RP-273).

Bankes, S. C. and Gillogly, J.: 1994a, "Validation of Exploratory Modeling", in High Performance Computing: Grand Challenges in Computer Simulation, A. Tentner (ed.), The Society for Computer Simulation, San Diego CA, pp. 382- 387 (also published as RAND RP-298).

Bankes, S. C. and Gillogly, J.: 1994b, "Exploratory Modeling: Search Through Spaces of Computational Experiments", in The Proceedings of the Third Annual Conference on Evolutionary Programming., A. V. Sebald and L. J. Fogel (eds.), World Scientific.

Bankes, S. C. and Gillogly, J., “A Software Environment for Exploratory Modeling: Support for Research Through Computational Experiments” (submitted to IEEE Software).

Darmstadter, J. and Toman, M. A. eds.: 1993, Assessing Surprises and Nonlinearities in Greenhouse Warming: Proceedings of an Interdisciplinary Workshop, Resources for the Future, Washing, DC.

De Geus, A. P.: 1988, “Planning as Learning,” Harvard Business Review, March-April, 70-74.

Dewar, J. A., Builder, C. H., Hix, W. M., and Levin, M. H.: 1993, Assumption-Based Planning: A Planning Tool for Very Uncertain

Times, RAND, Santa Monica, CA, MR-114-A, pp.
Dowlatabadi, H. and Morgan, M. G.: 1993, "Integrated Assessment of Climate Change," Science, 259, 1813.
Fickett, A. P., Gellings, C. W., and Lovins, A. B.: 1990, "Efficient Uses of Electricity," Scientific American, 263, 64–74.
Fisher, J. C., and Pry, R. H.: 1971, "A simple substitution model of technological change," Forecasting and Social Changes, 3, 75-88.
Hafele, W. (ed): 1981, Energy in a Finite World: A Global Systems Analysis. Balinger, Cambridge, Massachusetts, 253-278.
Hammitt, J. K., Lempert, R. J., and Schlesinger, M. E.: 1992, "A sequential-decision strategy for abating climate change," Nature, 357, 315–318.
Houghton, J. T., Jenkins, G. J. and Ephraums, J. J. (eds.): 1990, Climate Change The IPCC Scientific Assessment. Cambridge University Press, Cambridge, 364 pp.
Global Change Research Program: 1995, Our Changing Planet; The FY95 U.S. Global Change Research Program, A Report by the Subcommittee on Global Change Research, Committee on Environment and Natural Resources Research of the National Science and Technology Council, A Supplement to the President's Fiscal Year 1995 Budget.
Kauffman, S. A.: 1989, "Adaptation on Rugged Fitness Landscapes." In Lectures in the Sciences of Complexity, E. Stein, ed. Addison-Wesley, Reading, Mass.
Kolstad, C. D.: 1993, "Looking vs. Leaping: The Timing of CO2 Control in the Face of Uncertainty and Learning," in Costs, Impacts and Benefits of CO2 Mitigation, Proc. of Conference held at IIASA, Laxenburg, Austria, October 1992, Y. Kaya, N. Nakicenovic, W. Nordhaus and F. Toth (Eds.), International Institute for Applied Systems Analysis, pp. 63-82.
Kolstad, C. D.: 1994, "Mitigating Climate Change Impacts: The Conflicting Effects of Irreversibilities in CO2 Accumulation and Emission Control Investment," in Integrative Assessments of Mitigation, Impacts, and Adaptation to Climate Change, Proc. of Conference held at IIASA, Laxenburg, Austria, October 1993, N. Nakicenovic, W. D. Nordhaus, R. Richels and F. L. Toth (Eds.), International Institute for Applied Systems Analysis, pp. 205-218.
Lave L., and Dowlatabadi, H.: 1993, "Climate Change Policy: The effects of personal beliefs and scientific uncertainty," Environmental Science and Technology, 27, pp. 1692-1972.
Lempert, R. J., Schlesinger, M. E., Hammitt, J. K.: 1994, "The Impact of Potential Abrupt Climate Changes on Near-Term Policy Choices," Climatic Change, 26:351-376. Also published as: “The Impact of Potential Abrupt Climate Changes on Near-Term Policy Choices,” Lempert, R. J., Schlesinger, M. E., and Hammitt, J. K. in Integrative Assessment of Mitigation, Impacts, and Adaptation to Climate Change,

Proc. of Conference held at IIASA, Laxenburg, Austria, October 1994, N. Nakicenovic, W. D. Nordhaus, R. Richels and F. L. Toth (Eds.), International Institute for Applied Systems Analysis, pp. 173-204.

Lindzen, R. S.: 1990, "Some coolness about global warming," Bull. Amer. Meteorol. Soc., 71, 288-299.

Maier-Reimer, E., and Hasselmann, K.: 1987, "Transport and Storage of CO2 in the Ocean– An Inorganic Ocean-Circulation Carbon Cycle Model," Climate Dynamics, 2, 63-90.

Manne, A. S. and Richels, R. G.: 1991, "Global CO2 emission reductions: The impact of rising energy costs," The Energy Journal, 12, 88-107.

Manne, A. S. and Richels, R. G.: 1992, Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emissions Limits, MIT Press, Cambridge, 182 pp.

Mansfield, E.: 1961, "Technical change and the rate of imitation," Econometrica, 26, 741-766.

Morgan, M.G. and Herion, M.: 1990, Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis, Cambridge University Press, 332 pp.

National Center for Atmospheric Research (NCAR): 1994, El Niño and Climate Prediction, Report to the Nation on our Changing Planet, No. 3, Spring.

NAS: 1985, Policy Implications of Greenhouse Warming. National Academy of Sciences, Washington, DC, pp. 47-63.

Nordhaus, W. D.: 1991, "The Cost of Slowing Climate Change: A Survey," The Energy Journal, 12, 37-65.

Nordhaus, W. D.: 1992, "An optimal transition path for controlling greenhouse gases," Science, 258, 1315-1319.

Nordhaus, W. D.: 1994a, Managing the Global Commons: The Economics of Global Change, MIT Press, 213 pp.

Nordhaus, W. D.: 1994b, “Expert Opinion on Climate Change,” American Scientist, Vol. 82, January-February, 45-51.

OTA: 1991, Changing by Degrees: Steps to Reduce Greenhouse Gases. Office of Technology Assessment, U.S. Congress, OTA-0-482, Washington, DC, 354 pp.

Parson, E. A.: 1994. Searching for Integrated Assessment: A Preliminary Investigation of Methods and Projects in the Integrated Assessment of Global Climatic Change, paper prepared for the 3rd meeting of the CIESIN-Harvard Commission on Global Environmental Change Information Policy, NASA Headquarters, Washington DC, February 17-18.

Passell, P.: 1989, "Economic Watch; Curing the Greenhouse Effect Could Run Into the Trillions." New York Times, Section 1, Part 1, Page 1, Sunday, November 19.

Peck, S. C. and Teisberg, T. J.: 1992, "CETA: A Model for Carbon Emissions

Trajectory Assessment," Energy Journal, Vol. 13, No. 1 15, 55-77.
Peck, S. C. and Teisberg, T. J.: 1993, "Global Warming Uncertainties and the Value of Information: An Analysis Using CETA," Resource and Energy Economics 15, 71-97.
Rotmans, J.: 1990, IMAGE: Integrated Model to Assess the Greenhouse Effect, Kluwer, Boston, 289 pp.
Schelling, T.: 1983, "Climatic Change: Implications for Welfare and Policy", in Changing Climate, National Research Council.
Schlesinger, M. E., and Jiang, X.: 1990, "Simple model representation of atmosphere- ocean GCMs and estimation of the timescale of CO2-induced climate change," J. Climate, 3, 1297-1315.
Schlesinger, M. E., and Jiang, X.: 1991, "Revised projection of future greenhouse warming," Nature, 350, 219-221.
Schurr, S. H., Burwell, C. C., Devine, Jr., W. D., and Sonenblum, S.: 1990, Electricity in the American Economy: Agent of Technical Progress, Greenwood Press, Westport, CT, 443 pp.
Wack, P.: 1985a, “Scenarios: uncharted waters ahead,” Harvard Business Review, September-October, pp. 73-89.
Wack, P.: 1985b, “Scenarios: shooting the rapids,” Harvard Business Review, November-December, pp. 139-150.
Wagner, D.: 1992, Today then: America's best minds look 100 years into the future on the occasion of the 1893 World's Columbian Exposition. American & World Geographic Pub., Helena, MT, 226 pp.