

forthcoming in
Innovative Energy Strategies for CO2 Stabilization,
Cambridge University Press

Adaptive Strategies for Climate Change

Robert J. Lempert^a and Michael E. Schlesinger^b

a. RAND,
1700 Main St.
Santa Monica, CA 90407-2138.

b. Department of Atmospheric Sciences
University of Illinois at Urbana-Champaign
Urbana, IL 61801.

Table of Contents

1. Introduction.....	1
2. The Case for Adaptive-Decision Strategies.....	3
3. Assessing Adaptive-Decision Strategies.....	9
4. Design of Robust Adaptive-Decision Strategies.....	17
4.1 Impact of Variability on Adaptive-Decision Strategies.....	19
4.2. Promoting Innovation.....	29
5. Conclusions.....	42
References.....	47

1. Introduction

This book presents forecasts for a range of innovative energy technologies that might help stabilize atmospheric concentrations of greenhouse gases during the course of the 21st century. These forecasts provide a wealth of important information for those who wish to inform their view of the climate-change problem and the actions governments or the private sector might take to address it. But these chapters nonetheless present a fundamental dilemma -- for the one thing we know for sure about forecasts is that most of them are wrong. How then should we use the information in this book to shape policy?

The difficulty resides not so much in the forecasts themselves as in the methods that we commonly employ to bring the information they contain to bear on adjudicating among alternative policy choices. Generally, we argue about policy by first settling on our view of what will happen in the future and then by using this understanding to decide what actions we should take in response. For instance, if we came to believe, through arguments such as those in this book, that there were cost-effective technological means to stabilize greenhouse-gas emissions, we might be more likely to support policies that sought to achieve such a stabilization. We often make these arguments non-quantitatively, even if systematically. There are also a host of powerful mathematical tools, based on the mathematical techniques of optimization, that help us systematize and elaborate on this style of thinking about the future. These methods encapsulate our knowledge about the future in probability distributions which then allow us to rank the desirability of alternative strategies. These prediction-based analytic approaches work extraordinarily well in a wide variety of cases. So much so, that they strongly affect our notions of the criteria we ought to use to compare our policy choices and the way we ought to use forecasts to support these choices. We often speak of choosing the optimum or best strategy based on our predictions of the future. In this approach then, the purpose of forecasts is to shape our views of what is likely to happen, as a means of shaping our decisions about how to act.

In this chapter we argue for a different approach to climate-change policy and thus a different use for forecasts. We argue that climate change presents a problem of decision-making under conditions of deep uncertainty. We begin with the premise that while we know a great deal about the potential threat of climate change and the actions we might take to prevent it, we cannot now, nor are we likely for the foreseeable future, answer the most basic questions, such as is climate change a serious problem and how much would it cost to prevent it? We argue that in the face of this uncertainty, we should seek robust strategies. Robust strategies are ones that will work reasonably well no matter what the future holds. Not only is this desirable in its own right, but a robust strategy may provide a firm basis

for consensus among stakeholders with differing views about the climate-change problem because all can agree on the strategy without agreeing on what is likely to happen in the future. Rather than using forecasts to specify a particular path into the future, in this alternate approach forecasts describe a range of plausible scenarios and sharpen our sense of what any particular future, if it comes to pass, might look like.

In this chapter we argue that robust strategies for climate change are possible by means of adaptive-decision strategies, that is, strategies that evolve over time in response to observations of changes in the climate and economic systems. Viewing climate policy as an adaptive process provides an important reconfiguration of the climate-change policy problem. The long-term goal of the Framework Convention on Climate Change calls for society to stabilize atmospheric levels of greenhouse gases at some, currently unknown, safe level and the protocol negotiated in Kyoto in December 1997 commits the world's developed countries to specific, binding near-term emissions reductions. There is currently much debate as to how to best implement these reductions and whether or not they are justified. Viewed as the first steps of an adaptive-decision strategy, however, the Kyoto targets and timetables are but one step in a series of actions whose main purpose is to increase society's ability to implement large emissions reductions in the future. Contrary to most of today's assessments, the real measure of the Framework Convention's success a decade hence should not be any reductions in atmospheric concentrations of greenhouse gases, but rather the new potential for large-scale emissions reductions society has created for the years ahead.

In this context, the potential for new technologies, the processes of innovation and diffusion by which these technologies come into widespread use, and the government policies needed to encourage such processes, play a central role in society's response to climate change. We will suggest in this chapter that the processes of technology diffusion might provide key indicators for monitoring the progress of an adaptive climate-change strategy, particularly in the face of significant climate variability that will

make it difficult to observe a reliable signal about the extent of human-caused impacts to climate change. There is also much debate over the type of policy instrument -- carbon taxes or tradable permits -- that should be used to encourage emissions reductions. We find that one does not need to have particularly high expectations that these innovative technologies will achieve the potential described here in order to put in place an adaptive-decision strategy that also makes important use of a third type of policy instrument, technology incentives. Forecasts such as those in this book can provide a great deal of information that can help policy-makers design such robust adaptive strategies, but in order to be most useful, such forecasts need to provide a broader range and different types of information than is often provided.

We begin this chapter with a discussion of the importance of adaptive-decision strategies and how they differ from other views of the climate-change problem. We will then describe some first steps in employing an analytic method, based on a multi-scenario simulation technique called exploratory modeling, for designing and evaluating adaptive-decision strategies for climate change and provide a simple example of its application. We will then turn to the design of adaptive-decision strategies, examining the conditions under which such strategies are appropriate, the tradeoffs between actions and observations in the design of such strategies, and some initial steps in examining climate-change policy as a process. We will present two example analyses -- the impacts of climate variability on the design of such strategies and the choice of policy instruments in the presence of the potential for significant technology innovation. During the course of the discussion, we will show the role of innovation in emissions-reducing technologies in such adaptive-decision strategies, as well as the type of information from forecasts like those in this book that can prove most useful in designing them.

2. The Case for Adaptive-Decision Strategies

Much of the current policy and political debate about climate change focuses on the question of how much to reduce emissions of anthropogenic greenhouse gases. The Framework Convention on Climate Change, in particular the protocol negotiated in Kyoto in December 1997, now calls for the world's developed countries (Annex I nations in the protocol) to commit to binding reductions that would bring their emissions back to roughly 1990 levels by 2010. It does little damage to the nuances of the debate to divide most of the responses to this protocol into two opposing camps. Supporters argue that the risk of severe impacts from climate change is sufficiently high, and the prospects of inexpensive abatement due to technology innovation and social change sufficiently good, that immediate and significant reductions in emissions are the only prudent course. Opponents point to the potentially high economic costs of emissions reductions and the large uncertainties surrounding the predicted damages, and argue that little or no action be taken at this time.¹

Numerous quantitative policy studies have attempted to determine the relative validity of these competing arguments. Many take the Kyoto targets as a given and attempt to predict the economic cost of meeting them. These estimates cover a wide range, depending on assumptions about diverse factors, such as how the protocol is implemented and the market acceptance of efficiency-enhancing and emissions-reducing technologies (Repetto and Austin, 1997). In addition, many analytic efforts have been made to determine the best arrangements for the emissions trading and other policies that may be used to implement the protocol. For instance, economic studies generally argue that allowing countries to trade emissions permits among themselves will lower the costs relative to requiring each country to meet its reduction target individually.

More broadly, other studies address the question of whether or not the Kyoto targets represent the most cost-effective first steps towards the Climate Convention's ultimate goal of long-term stabilization of atmospheric concentrations of greenhouse gases. The influential study by Wigley, Richels, and Edmonds (Wigley et al., 1996 -- hereafter WRE) argued that the

Convention's goals can be met at least cost with near-term emissions reductions much smaller than those required by Kyoto. In response, a number of studies have questioned the assumptions underlying WRE. In particular, these rebuttals argue that alternative, plausible assumptions about the speed and characteristics of technology innovation suggest that, given the large time scales associated with changing energy-related capital infrastructure and the feedback effects whereby early investments in new technology can reduce subsequent costs, Kyoto's emissions reductions would be beneficial.

It is not at all clear that the question of the best level of near-term emissions reductions can be resolved by this type of argument. Ultimately, the predicted benefit of such reductions depends on a broad range of assumptions about the seriousness of potential impacts due to climate change, the potential of new emissions-reducing technologies, the value people in the future will place on the preservation of nature, the extent to which market reforms succeed in many developing countries, and a host of other factors that are impossible to predict with any accuracy. Thus, the studies that favor and oppose particular levels of emissions reductions often differ, not so much because of any disparities in the underlying methodology, but in the basic assumptions they make about the future. Because a wide range of such assumptions is plausible, the studies cannot distinguish among a wide range of potential policy proposals. Furthermore, such studies will provide little basis for agreement among competing sides in the climate-change-policy debate, since each stakeholder will, not surprising, tend to gravitate towards those assumptions that justify the level of emissions-reductions they would otherwise support on ideological, financial, or other grounds.

The concept of adaptive-decision strategies offers a way out of this impasse. Rather than base prescriptions for the proper level of emissions reduction on a judgment that WRE is or is not correct, or that some new emissions-reducing technology will or will not achieve substantial market penetration, we look to strategies that begin with a certain suite of actions,

monitor the external environment, and then modify the suite in response to these observations. Rather than predict what the level of emissions will be in future decades, we determine what type of adaptive-decision strategy will perform best across a wide range of potential futures, each of which in retrospect, clearly requires some different level of emissions reduction.

Simple and obvious as this point may seem at first, it is actually a substantial reconfiguration of the policy problem because it suggests that policy-makers consider a wider set of potential policy actions and a different set of information about the world than they are encouraged to do in the standard analytic framework. Perhaps this can best be seen by comparing the concept of adaptive-decision strategies to the related, analytic approach of sequential-decision strategies which more commonly appears in the climate-change-policy literature (Manne and Richels, 1992; Nordhaus, 1994; Hammitt, Lempert, and Schlesinger, 1992 -- hereafter HLS). The sequential-decision approach finds an optimum strategy assuming that policy-makers begin with uncertainty but will receive a specific dollop of uncertainty-reducing information at some fixed time in the future. For instance, HLS consider two-period decision strategies that start with imperfect information about the future costs of climate change and the costs of abatement, with the uncertainty resolved completely in ten years. Generally such a study suggests some modest adjustment to the strategy one would pursue without uncertainty, such as a slightly increased level of near-term emissions reductions.

A sequential-decision strategy assumes that the rate at which policy-makers learn is independent of the seriousness of the problem or of actions they take. But the dependence of what we learn tomorrow on both what we do today and the possible paths the world offers to us may be important factors in choosing a response strategy. For instance, a key determinant of the proper level of near-term emissions reductions ought to be whether or not society can achieve sufficient warning time to respond to signals of adverse climate change. With sufficient warning time, near-term emissions can be lower; without it they should be higher. If climate change is modest, it may

take decades to extract the signal from the noise (Kelly, Kolstad, and Schlesinger, 2000). But if climate change is severe, the signal may well be revealed more quickly. Similarly, near-term emission reductions will likely generate information about abatement costs that can help us better determine the proper level of emissions reductions. Including such considerations in the calculations will affect analysts' recommendations as to the best level of near-term emissions reductions.

Considering strategy as an evolutionary process can also significantly affect the types of actions we consider and the criteria we use to judge those actions. For instance, a key question in the design of adaptive-decision strategies is the information to which decision-makers ought to pay attention. For example, the U.S. Federal Reserve follows an adaptive-decision strategy as it periodically adjusts interest rates to regulate the growth of the U.S. economy. Much attention is paid to the information and indicators the Fed governors use to make their decisions and, when the Fed chairman makes comments that additional information will be deemed important, the markets can react strongly. In fact, sometimes indicators alone can be a large step towards the solution. One of the most striking examples was the U.S. EPA's policy of requiring companies to publish their toxic-release inventories. This communication of information itself reduced emissions, without any other regulation, because companies felt shamed by high numbers.

Perhaps most importantly, the response of real-world decision-makers to deep uncertainty is often not any particular action, but rather on building processes that can respond to uncertainty. This in the broadest sense is what we mean by an adaptive-decision strategy. In his review of risk assessment, William Clark (1980) emphasizes the importance of rules of evidence and other such processes in determining how we perceive and response to risk. The business literature often recommends organizational changes, designed to increase the firm's ability to respond to new information, as the best response to uncertainty and rapid change. This theme of information flows and how to use them is central to the idea of an adaptive-decision strategy. It may be that the most important responses to the threat of climate change are

the establishment of the institutional mechanisms that allow society to respond more quickly and less expensively to opportunities and threats in the decades ahead.

3. Assessing Adaptive-Decision Strategies

Despite the fact that adaptive-decision strategies are likely to be the better approach to climate change, and the approach policy-makers are likely to follow in practice, they have been infrequently considered in the policy literature (see Dowlatabadi, 1999, for one of the few such treatments). One important reason is that the analytic-policy literature has been largely dominated by the approach of finding the optimum response to climate change. This optimization approach is not conducive to assessing adaptive-decision strategies for two reasons. First, the analytic demands of optimization calculations usually require neglecting key feedbacks. Often these feedbacks are precisely those involved in an adaptive-decision strategy.

More profoundly, optimization is the wrong criterion for assessing adaptive-decision strategies because it assumes a level of certainty about the future that, if truly present, would obviate most if not all of the need for such strategies. Optimization finds a unique, best strategy based on specific assumptions about the system model, the prior probability distributions on the parameters of that model, and a loss function which represents society's values. But adaptive-decision strategies are most useful when there is deep uncertainty, that is, when we do not know with confidence the model, probabilities, or societal values, or where different stakeholders disagree about these things. In such cases, disagreements about optimum strategies can quickly reduce to arguments about alternative, unprovable assumptions or differences in goals.

Rather than optimization, the criterion for assessing adaptive-decision strategies ought to be robustness. Robust strategies are ones that will work reasonably well, at least compared to the alternatives, over a wide range of plausible futures. Robust strategies are advantageous because we can often identify them without specifying, or requiring stakeholders to agree on,

specific models, priors, or societal values. In general, we can always identify, post hoc, models, priors, and values that make any given adaptive-decision strategy optimal. But in practice, the robustness criterion is useful precisely because it avoids the need for any prior agreement in those cases where a range of plausible scenarios is the best available representation of the information we have about the future.

The concept of robust strategies has a strong theoretical and practical pedigree. The idea is closely related to Simon's (1959) satisficing strategies and Savage's (1954) idea of minimizing the maximum regret. There is also much evidence that actual decision-makers in practice search for robust strategies rather than optimal ones (March, 1994). While these ideas have been familiar in the decision-analysis literature (Watson and Bruede, 1987; Matheson and Howard, 1968), in practice they are infrequently employed on the climate-change problem because they are often difficult to implement for problems of any complexity. In recent years, however, an emerging school of what we might call multi-scenario simulation approaches (Lempert, Schlesinger, and Bankes, 1996, henceforth LSB; Morgan and Dowlatabadi, 1996; Rotmans and de Vries, 1997; Casan, Morgan, and Dowlatabadi, 1999), has begun to exploit the capabilities of new computer technology – fast computer processing; extensive, low-cost memory; and powerful, interactive visualization tools – to consider strategies under deep uncertainty. These methods use simulation models to construct different scenarios and, rather than aggregate the results using a probabilistic weighting, instead make arguments from comparisons of fundamentally different, alternative cases.

Our approach, which we call exploratory modeling (Bankes, 1993), is a multi-scenario simulation technique that explicitly implements these classic ideas about robust strategies. In exploratory modeling we use simulation models to create a large ensemble of plausible future scenarios, where each member of the ensemble represents one guess about how the world works and one choice among many alternative strategies we might adopt to influence the world. We then use search and visualization techniques to extract from this ensemble of scenarios information that is useful in distinguishing among

policy choices. These methods are consistent with the traditional, probability-based approaches to uncertainty analysis because when such distributions are available, one can lay them across the scenarios and thus calculate expected values for various strategies, value of information, and the like. However, in situations characterized by deep uncertainty, the exploratory modeling method allows us to systematically find strategies that are robust against a wide range of expectations about the future and a range of valuations of that future.

We now provide a simple example that demonstrates how this concept of robustness and these exploratory modeling methods can be used to assess adaptive-decision strategies. This example, based on our work in LSB addresses whether and under what conditions adaptive-decision strategies are a reasonable response to climate change. We begin by comparing the performance of the strategies advocated by the two competing camps in the debate over the Kyoto protocol. We define two alternative strategies, "Do-a-Little" (DAL) and "Emissions-Stabilization" (ES), to represent these camps. Each strategy reflects a prediction-based approach to climate change and is expressed as a given emissions path over the course of the 21st century. The former has little near-term reductions and is similar to the results of many economic analysis that assume relatively small long-term damages due to climate change (less than 2% of GWP) and relatively high abatement costs (more than 2% of GWP) to stabilize atmospheric concentrations of greenhouse gases. The latter, which returns global emissions to their 1990 levels by 2010 and holds them there until mid-century, is similar (though slightly more aggressive) than the reductions paths mandated by the Kyoto protocol. While the proponents of these camps do not explicitly reject adaptive-decision approaches, and the Framework Convention on Climate Change explicitly calls for policies that adjust over time in response to new information, the debate in practice is very much characterized by support or opposition to specific targets for the reductions of greenhouse gases.

We compare these DAL and ES strategies using a linked system of simple climate and economic models. 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. We work in a cost-benefit framework, and measure the performance of each strategy as the present value of the sum of the hundred-year time series of damage and abatement costs. In particular, we focus on comparing the performance of these strategies as a function of three key dimensions of uncertainties we face about climate change (Lave and Dowlatabadi, 1993; Nordhaus, 1994), which we express as: (i) the sensitivity of the climate system to increasing concentrations of greenhouse gases, (ii) damages resulting from an increase in global-mean surface temperature, and (iii) the ability of innovation to significantly reduce the costs of abating greenhouse-gas emissions. In each case we define our plausible range of estimates, wherever possible, as the extreme range found in the published, refereed literature (a similar screening was used in Rotmans and de Vries, 1997).

We simulate the impact of uncertainty about the climate sensitivity on the change in global-mean surface temperature due to anthropogenic emissions of greenhouse gases (GHGs) with our energy-balance-climate/upwelling-diffusion-ocean (EBC/UDO) model (Schlesinger and Jiang, 1991; Schlesinger and Ramankutty, 1992, 1994a,b, 1995; and Schlesinger, 1998). We allow the climate sensitivity to vary between 0.5°C and 4.5°C. The upper value is the high value from the IPCC; the low value is from Lindzen (1990). We express uncertainty about the damages due to climate change using a simple phenomenological impacts function of changes in global-average temperature, as in the practice of many integrated assessments (Nordhaus, 1994a; Manne and Richels, 1992; Peck and Teisberg, 1993, 1992). We consider cases where total aggregate damages at the end of the next century for a 3°C temperature rise range from 0.5% to 20% of gross world product (GWP), based on a survey of experts conducted by Nordhaus (1994b).

The crude innovation model used in LSB to simulate the consequences of uncertainty about the impacts of innovation, first used in HLS, assumes basecase greenhouse-gas emissions (Houghton et. al., 1990) are reduced as non-emitting "fuel-switching" technologies diffuse through the economy with an S-shaped, logistic diffusion curve at some policy-determined rate $1/R$. (We will describe a more detailed innovation model below). The model builds the least-cost mix of emitting and non-emitting technologies to meet the exogenous energy-demand and policy-imposed emissions constraint. The model parameters are chosen so that our basecase innovation case reproduces the results of more detailed models (Manne and Richels, 1991; Nordhaus, 1991). We then assume that innovation can reduce the incremental costs of the non-emitting technologies at some fixed, but currently unknown, annual rate. Abatement costs also depend on the rate of reductions because, in those cases where emissions are reduced sufficiently quickly, emitting capital must be prematurely retired (assuming a normal lifetime of 30 years). This formulation captures in a crude manner the idea of inertia that may affect the choice between early and late action (Grubb, Chapuis, and Ha-Duong, 1995).

Figure 1 compares the performance of the DAL and ES strategies as a function of society's expectations about our three, key climate-change uncertainties: the likelihood of extreme climate sensitivity, the likelihood of extreme damages, and the likelihood of significant innovation. For instance, the lower left-hand corner represents the expectation that the climate is likely to be insensitive to increasing greenhouse-gas concentrations, damages due to any climate change are likely to be small, and that innovation is unlikely to reduce the costs of abatement. In contrast, the upper right-hand corner represents the expectation that the climate system is very sensitive to greenhouse gases, damages are likely to be large, and innovation is likely to radically change the costs of abatement. The curved surface labeled A represents the boundary where we should be indifferent between the two choices, DAL and ES. Formally, this is the indifference surface where the strategies have equal expected values. To the left of the curve, we should prefer DAL; to the right we should prefer ES.

We created this visualization by running a large number of scenarios, each with one of the two strategies and a particular set of assumptions about the three parameters describing our uncertainties. We then laid probability distributions across these parameters as a function of these three expectations, and conducted a computer search to find the surface on which the expected values of the two strategies were equal. Note that this process, similar to the policy-region analysis of Watson and Bruede (1987), inverts the standard use of probabilities in decision analysis. Commonly, analysts begin with assumptions about the probability distributions of key uncertainties, which much be gleaned from methods such as expert elicitation (Morgan and Henrion, 1990), and then use these probabilities to find some optimum strategy. Instead, we report key probabilities as outputs of the calculations. That is, we report the probabilities that are implicit in advocating one strategy over another.

There are two important arguments we can make based on the comparison shown in Figure 1. First, the figure provides what we call a *landscape of plausible futures* (Park and Lempert, 1998), because it provides an explicit overview of the implications of the full range of uncertainties. Such landscapes are useful because they provide a common framework in which contending stakeholders to the climate-change debate can position themselves. Each stakeholder can find a portion of the landscape that reflects their view of the world and thus agree that the models we are using capture their assumptions and can reproduce the arguments they make from those assumptions. This process helps stakeholders find a language in which to express their disagreements. It also helps them buy-in to the analysis. We have shown this "climate cube" to audiences including both oil-company executives and ardent environmentalists and convinced them that their divergent views are captured within different regions in the cube.

Second, this figure shows why DAL and ES provide an unsatisfactory set of options. We cannot predict which point in this cube best represents the state of the world. Little scientific evidence is currently available, or is likely to be available anytime soon, to convince a reasonable individual whose beliefs place her on one side of surface A that she in fact should switch to the

other camp. For instance, someone who believes the new, emissions-reducing technologies are likely to make future greenhouse-gas emission reductions inexpensive can offer dozens of examples of technologies that have in the past, and others that may in the future, dramatically change the cost of emissions reduction and other pollution prevention. Alternatively, someone who believes such emissions-reducing technologies are unlikely can point to dozens of stories about over-optimistic forecasts and technologies that failed to change the world. Because the technological future is fundamentally unpredictable, there is no evidence we can gather nor theorems we can evoke that will definitely sway this argument one way or the other. Thus, we cannot really know which side of the surface A in Figure 1 we are on. We can show, however, that guessing wrong can be very costly. That is, if we chose ES in a world where it turns out DAL would have been best, or DAL in an ES world, the costs can be significant. Thus, it should be no surprise that framing the climate-change problem around competing predictions of an unpredictable future should foster a hostile and unresolvable debate.

We next consider how an adaptive-decision strategy performs in comparison to these two static strategies. We posit a very simple, threshold-based adaptive-decision strategy that observes the damage and abatement costs once a decade and can make one correction based on these observations. The strategy begins with slow emissions reductions, as shown in Figure 2. If the damage exceeds a particular threshold after ten years, the strategy switches to draconian emission reductions. If not, the strategy waits another ten years. If the damage then exceeds the threshold, or if the abatement costs are below a threshold, then the strategy switches to rapid reductions. If not, it continues checking the damage and abatement costs every decade until the mid-twenty first century. We assume that once the strategy makes a mid-course correction, it makes no further observations nor reduction-rate adjustments. This particular adaptive-decision strategy is one, very simple exemplar, not necessarily the best such strategy, but sufficient for our purposes here.

Note that we express this strategy, not as any particular outcome for emissions-reductions, but as a specific process that will produce different

levels of emissions reductions in alternative scenarios. Thus, in a crude way, we can examine the type of process the decision-makers ought to use rather than the particular outcomes they ought to achieve. This question is different from those generally addressed by optimization approaches which, by definition, assume that the decision-maker conducts a particular, idealized process. Our claim, however, is that modeling the process of decision-making not only helps answer questions more relevant and credible to decision-makers faced with deep uncertainty, it also (and probably not coincidentally) provides solutions that perform better across a wide range of scenarios than those based on optimum outcomes.

We compare the performance of this adaptive-decision strategy to that of the DAL and ES strategies, with the curves labeled B and C in Figure 1. Surface B, in the lower left-hand corner of the cube, is the indifference surface between the adaptive-decision strategy and DAL. If society's expectations about the future put us in this lower left-hand region, we should prefer DAL; otherwise we should prefer adaptive-decision. Surface C, in the corner of the cube closest to the viewer, is the indifference surface between the adaptive-decision strategy and ES. If society's expectations about the future put us in this upper-right-hand region, then we should prefer ES; otherwise we should prefer adaptive-decision. In addition, the adaptive-decision strategy never makes large errors, so the costs of choosing it incorrectly -- where DAL or ES in fact turn out to be the right answer -- are not large.

The adaptive-decision strategy performs better than DAL and ES over a wide range of plausible expectations about the future, and thus provides a more robust response than either of the two static strategies. For those expectations where the adaptive-decision strategy does not perform better than DAL or ES, the cost of choosing it compared to the best option is not large. (Later in this chapter we will formalize this notion of robustness.) In addition, if we convinced a group of opposing stakeholders that the different poles of the climate-change debate are captured within this cube, then they should now agree that the adaptive-decision strategy is the proper way to

address the climate-change problem, whether or not they previously resided in the DAL or ES camp.

4. Design of Robust Adaptive-Decision Strategies

It is perhaps not surprising that an adaptive-decision strategy which can evolve over time in response to new information is more robust against a wide range of uncertainty than strategies which do not adapt. In order to shape policy choices, however, we need to determine the best adaptive-decision strategies from among the options available. In the above example, we examined only a single, threshold-based, adaptive-decision strategy, which started with slow near-term emissions reductions and used particular trigger levels for its observations of damage and abatement costs. In order to design and use adaptive-decision strategies, we need to answer question such as: should such strategies begin with fast or slow emission reductions? What observations should suggest a change in the initial emissions-reduction rate?

Several academic literatures offer insights into the design of adaptive-decision strategies. Economics and financial theory suggest that a successful strategy will employ a portfolio of different policy instruments and the mix and intensity of these instruments may change over time. Control theory suggests that there is an intimate relation between our choice of instruments and the types and accuracy of the observations we can make. Scenario planning offers a useful language for the components of an adaptive-decision strategy: *shaping actions* intended to influence the future that comes to pass, *hedging actions* intended to reduce vulnerability if adverse futures come to pass, and *signposts* which are observations that warn of the need to change strategies (Dewar et al., 1993, van der Heijden, 1996).

We can use our exploratory-modeling methods to systematically combine these seemingly disparate elements and find robust adaptive-decision strategies. Generalizing on the application described above, we compare the performance of a large number of alternative adaptive-decision strategies across the landscape of plausible futures. In order to make sense of this information, we calculate the regret of each strategy in each of many plausible states of the world, defined as the difference between the

performance, for that state of the world, of the strategy and of the optimal strategy. We then search for robust strategies, defined as those with relatively small regret over a wide range of states of the world. In some cases we can find strategies that are robust across the landscape of plausible futures. In other cases, we find particular states of the world that strongly influence the performance of robust strategies. We then might report tradeoff curves which show, for instance, the choice of robust strategy as a function of the probabilistic weight a decision-maker might assign to some critical state of the world.

In this section we will use these methods to address two important questions in the design of adaptive-decision strategies. First, we will examine the impacts of climate variability on the choice of near-term policy choices (Lempert, Schlesinger, Bankes, and Andronova, (2000) – henceforth LSBA). Variability, one of the most salient features of the earth's climate, can strongly affect the success of adaptive-decision strategies by masking adverse trends or fooling society into taking too strong an action. Interestingly for this book, our preliminary analysis suggests that the most robust adaptive strategies in the face of climate variability begin with moderate near-term emission reductions which are more likely to shift to more rapid abatement based on observations of innovation-reduced abatement costs than of increased climate damages. Second, we examine the mix of policy instruments an adaptive-decision strategy ought to use to encourage the diffusion of emissions-reducing technologies [Robalino and Lempert (2000); henceforth RL]. We find that in many circumstances a mix of price-based mechanisms (e.g., carbon taxes or tradable permits) and direct technology incentives are more robust than a single price-based mechanism alone. This analysis provides some direct lessons for the type of information that technology forecasts can most usefully provide.

4.1 Impact of Variability on Adaptive-Decision Strategies

In addressing the impacts of climate variability on adaptive-decision strategies, it is important to first note that an adaptive strategy is not always the best strategy to follow if you take into account the costs of adapting. For

instance, there might be expensive monitoring equipment needed to gather information, adjustment costs every time a strategy is changed, and/or the observations used to inform the adaptive strategy might be ambiguous so that mistakes are possible. If the costs of adapting are greater than the expected benefits, it is best to just ignore any signposts and make a permanent, for the foreseeable future at least, choice of shaping and hedging actions. This is the message in the lower-left and upper-right hand corners of Figure 1. In these regions the costs of waiting to observe new information outweigh the expected benefits of acting on that information.

The risk of adapting incorrectly is a key cost of an adaptive-decision strategy. Observations are often ambiguous, especially if the observed system has noise and other fluctuations. Decision-makers must balance between waiting until the information becomes more clear and the risk of acting on erroneous information (often called Type I and Type II errors). Given the degree of variability in the climate system, these dangers may be acute for climate-change policy. We need to ask if adaptive-decision strategies can still be robust in the face of climate variability, and if so, what these strategies are.

Our work in LSBA is among the first studies of the impacts of variability on climate-policy choices. While a simple, preliminary treatment, it nonetheless provides some important insight into the design of adaptive-decision strategies. We begin by making two changes to the models used in LSB in order to represent climate variability and its impacts. First, we treat climate variability as a white-noise component to the radiative forcing, so that the change in forcing is given by

$$\Delta Q(t) = 6.3334 \ln\left(\frac{\text{ECD}(t)}{\text{ECD}(1765)}\right) + \Delta F_{\text{SO}_4} \left(\frac{\text{E}_{\text{SO}_2}(t)}{\text{E}_{\text{SO}_2}(1990)}\right) + g(t) \quad , \quad (1)$$

where $\text{ECD}(t)$ is the effective carbon dioxide concentration that would give the same radiative forcing as the actual concentration of carbon dioxide, methane, and other greenhouse gases; $\text{E}_{\text{SO}_2}(t)$ is the emission rate of SO_2

which is converted to sulfate (SO_4) aerosols in the atmosphere; ΔF_{SO_4} is the sulfate-aerosol radiative forcing in 1990; and $g(t)$ is Gaussian-distributed noise with mean zero and standard deviation σ_Q (e.g., Hasselmann, 1976) which generates a red-noise-like temperature variability.

We represent our uncertainty about the climate system with a range of plausible values for the climate sensitivity, ΔT_{2x} , and the sulfate forcing, ΔF_{SO_4} . We choose the range of climate sensitivities, $0.5^\circ\text{C} \leq \Delta T_{2x} \leq 4.5^\circ\text{C}$, as in LSB. In practice we can find a best-guess estimate of the climate-sensitivity-dependent sulfate forcing parameter by using the EBC/UDO model to reproduce the instrumental temperature record (Schlesinger and Andronova, 2000). However, there are many sources of error, such as the influence of volcanoes or solar-irradiance variations, that could affect these estimates. Thus, we characterize our uncertainty about the sulfate forcing by the percentage it deviates from the best-estimate value as a function of climate sensitivity. For any pair of values for the sulfate forcing parameter and the climate sensitivity, we find the best estimate of the associated white-noise climate forcing, σ_Q , by regressing our EBC/UDO climate model using each pair against the instrumental temperature record from 1856 to 1995.

We represent our uncertainty about the impacts of climate change using a simple, phenomenological damage function designed to capture, in aggregate, some of the impacts of climate variability and the ability of society to adapt to changes in variability. We write the annual damage in year t , as percentage of Gross World Product (GWP), as

$$D(t) = \alpha_1 \left[\frac{\Delta \bar{T}_5(t)}{3^\circ\text{C}} \right]^{\eta_1} + \alpha_2 \left[\frac{\Delta T(t) - \Delta \bar{T}_5(t)}{0.15^\circ\text{C}} \right]^{\eta_2} + \alpha_3 \left[\frac{\Delta T(t) - \Delta \bar{T}_{30}(t)}{0.35^\circ\text{C}} \right]^{\eta_3}, \quad (2)$$

where $\Delta T(t)$ is the annual global-mean surface temperature change relative to its pre-industrial value, and $\Delta \bar{T}_5(t)$ and $\Delta \bar{T}_{30}(t)$ are the 5-year and 30-year running averages of $\Delta T(t)$.² The second term on the right-hand side of Eq. (2) represents impacts due to changes in the variability of the climate

system that society and ecosystems can adapt to on the time-scale of a year or two. The third term represents those impacts that society and ecosystems adapt to more slowly, on the order of a few decades, and thus are sensitive to both the year-to-year variability and the secularly increasing trend in temperature. The first term represents the damages due to a change in the global-mean surface temperature and is similar to the power-law functions used in most simple damage models in the literature, which we can view as representing impacts that society and ecosystems adapt to on century-long time scales. The coefficients α_1 , α_2 , and α_3 represent, respectively, the damages due to a 3°C increase in the global-mean temperature and the maximum damages due to climate variability in 1995 at the 90% confidence level.

We use a variety of empirical data to fix the damage-function parameters in Eq. (2). Because there are many gaps in the available information, and much of it that exists is heavily debated, we define a wide range of plausible parameter combinations rather than any best estimate. In LSBA we used only the requirement that the model be consistent with past observations to constrain the input parameters, though in general, we could also use future forecasts to generate constraints as well (as in the work on innovation described below). The wide range of plausible parameters supports, rather than hinders, our goals in this study since in the end we show that a simple adaptive-decision strategy can be robust against both very small and very large damages.

We constrain the parameters for the first term in Eq. (2) by noting that whatever damages due to climate change have occurred in the last few years and decades, they cannot have been more than a few tenths of a percent of GWP, otherwise we would have observed unambiguous evidence of damages to date. With $\Delta\bar{T}_5(1995) = 0.5^\circ\text{C}$, we can write $\alpha_1 \leq 0.1\% \cdot 6^{n_1}$, which corresponds to a range for the damage coefficient $\alpha_1 \leq 20\%$ GWP for cubic damages. We constrained the parameters in the second and third terms using time series data on economic losses from large-scale natural disasters (Munich Re Reinsurance, 1997). This is an imperfect data source because these extreme-

event damages (which in 1996, excluding earthquakes, caused \$60 billion, 0.2% GWP, in damages) are due at least in part to trends in society's vulnerability to natural disasters rather than to any change in the size or frequency of natural disasters themselves and, conversely, are only one component of the damages due to climate change. Lacking better sources of available information, we used these extreme-event data to define three bounding cases: "Low Variability," with parameters $(\alpha_2, \alpha_3, \eta_2, \eta_3) = (0.2\%, 0\%, 1, na)$; "High Variability" with parameters $(0.4\%, 0\%, 2, na)$; and "Increasing Variability" with parameters $(0\%, 0.33\%, na, 3)$. (Note that the "Low" and "High" damages due to variability cases drop the third term in Eq. (2), while the "Increasing Variability" case drops the second.) The "High" and "Increasing" cases differ in that damages in the latter grow with increasing concentrations of greenhouse gases while the former do not. Thus damages due to climate variability can be affected by emissions-abatement policy in the latter case but not the former.

This damage model has important shortcomings. Among the most important, the white-noise forcing, the driver of the variability in our model, is fit to the instrumental temperature record and does not change as we run our simulations into the future. Thus the damage distribution due the variability terms change over time only due to increases in the rate of change in the 30-year average temperature, though we expect that in actuality changes in variability are more important than changes in the mean (Katz and Brown, 1992; Mendelsohn et. al., 1994). Nonetheless, the crude phenomenological damage function in Eq. (2) provides a sufficient foundation to support initial explorations of alternative abatement strategies and the impacts of variability on near-term policy choices.

We can now use this model to examine the potential impacts of climate variability on near-term policy choices. Figure 3 shows the difficulties an adaptive-decision strategy might have in making observations of the damages due to climate change. The thin lines show the damages due to climate change generated by Eq. (2) for two distinct cases. In the first there is a large trend but low variability. In the second there is a large variability, but no trend.

The thick solid line in each case shows the trend a decision-maker (that is, a decision-making process that includes the making and reporting of scientific observations) might reasonably infer from the respective damage time series, calculated using a linear Bayesian estimator that rapidly detects any statistically significant trend in the damage time series. The estimator does a reasonably good job of tracking the damage time series in both cases, but because of the high variability, the estimates for the trend and no-trend cases do not diverge until about 2020. Thus, an adaptive-decision strategy attempting to distinguish between these two cases based on observations of the damage time series would have to wait at least two decades before being able to act. While there are many cases that the estimator can distinguish more quickly, the question nonetheless remains, how can we design an adaptive-decision strategy that can perform successfully when the observations have the potential for such ambiguity?

In order to address this question, we posit a simple, two-period, threshold-based adaptive-decision strategy, similar to the one used in LSB. The strategy can respond to policy-makers' estimate of any trend in damages based on annual observations of the noisy time series $D(t)$ or to any observed changes, neglecting noise, in abatement costs, represented here by the incremental cost of non-emitting capital. As shown in Figure 4, our adaptive strategies begin with a pre-determined abatement rate $1/R_1$ and can switch to a second-period abatement rate $1/R_2$ in the year, t_{trig} , when either the damages exceed, or the abatement costs drop below, some specified target values. The logistic half-life R represents the years needed to reduce emissions to one-half the basecase. The damage target (in % GWP) is given by $D_{\text{est}}(t) > D_{\text{thres}}$ ($D_{\text{est}}(t)$ is the decision-makers' estimate shown in Figure 3) and the abatement-cost target (in \$/ton carbon abated) is given by $K(t) < K_{\text{thres}}$. The second-period rate depends on the year t_{trig} . If $t_{\text{trig}} < T_{\text{near}}$, then the second-period abatement is given by $R_2 = R_2^{\text{near}}$. If $T_{\text{near}} \leq t_{\text{trig}} < T_{\text{far}}$, the second-period abatement is $R_2 = R_2^{\text{mid}}$. If neither condition is met by the year T_{far} , the strategy switches to a second-period

abatement $R_2 = R_2^{\text{far}}$. We express the decision facing policy-makers as a choice among the eight parameters defining these adaptive-decision strategies.

We now compare the performance of a large number of alternative adaptive-decision strategies. Figure 5a compares the expected regret of alternative strategies to the expected regret of the static DAL and ES strategies as a function of our expectations about the future. The horizontal axis gives the likelihood we ascribe to the possibility that DAL, as opposed to ES, is the better response to the climate-change problem. On the left-hand-side of the graph we are sure that DAL is better. Not surprisingly, the expected regret of the DAL strategy at this point is small (\$0.2 billion/ year) while the regret of the ES strategy is high (\$91 billion/ year). On the right-hand-side we are sure that ES is better, and not surprisingly the expected regret of ES is small (\$0.9 billion/year) here and the expected DAL regret is large (\$230 billion/year). In the middle of the graph, we ascribe even odds that DAL or ES is the best strategy.³

This figure is similar to the cube in Figure 1, except that we have collapsed our expectations about the future -- the independent variables -- into a single dimension so that we can compare the relative cost of the strategies. There is another major difference. In Figure 1 we considered only one adaptive-decision strategy; here we examine the performance of thousands (5,120 to be precise). Note that we have found adaptive-decision strategies that dominate the DAL and ES strategies no matter what our expectations about the future. That is, this chart shows that even with the potential of potentially ambiguous observations, adaptive-decision strategies are still the best response to climate change.

Of course, not all possible adaptive-decision strategies are better than ES and DAL. In fact, Figure 5a shows only a very small number of dominant strategies, those that perform best compared to the others as a function of expectations about the future. For each of many points across the horizontal axis we find that strategy with the lowest expected regret and label each of these dominant strategies with the three parameters -- the rate of near-term emissions reductions, the damage threshold, and the innovation threshold --

$(R_1, D_{\text{thres}}, K_{\text{thres}})$, that are most relevant to policy-maker's near-term decisions. As our expectations increase that ES, as opposed to DAL, is the better strategy, we pass through five different adaptive-decision strategies, with R_1 ranging from infinity to 40 years.

Two interesting patterns emerge as one examines this set of best adaptive strategies as a function of expectations about the future. First, there is a tradeoff between the rate of near-term emissions reductions and the confidence one should require in observations of damage trends before acting on them. That is, in the face of variability in the climate system, policy-makers can choose a response threshold for observed damages that can compensate, to a greater or lesser extent, for any choice of near-term emissions-reduction target. Second, the rate of near-term emissions reductions depends most sensitively on expectations about the future, while the aggressiveness with which society ought to respond to innovations are the least sensitive. That is, the policy choice at the focal point of the current negotiations may be the most controversial component of a robust strategy, in part because stakeholders with different expectations will have the most divergent views as to the proper target, while the least controversial components may be at the periphery of the negotiations.

Figure 5a also shows the most robust strategy across the range of expectations, which we find by searching for the strategy with the least squares regret across the points on the horizontal axis. This most robust strategy, given by (60 yrs, 1.2%, \$65), begins with moderate emissions reductions, is relatively insensitive to observations of climate damages, and is very sensitive to observations of decreasing abatement costs. Its annualized, present-value expected regret is \$22 billion/year, independent of our expectations about the ES future. While this low number does not mean climate change is costless (the regret measures the difference between the performance of a given strategy and that of the best strategy for that state of the world), it does suggest that even faced with deep uncertainty, we can find adaptive-decision strategies that perform nearly as well as the strategies we would choose if we knew the future with clarity.

The adaptive-decision strategies shown in Figure 5a still contend with a rather narrow range of uncertainty. For instance, in making this figure, we have assumed equal likelihoods of each of our variability cases -- Low, High, and Increasing -- that sulfate emissions are low, as in the scenarios of the Special Report on Emissions Scenarios (SRES; Nakicenovic et al., 2000), and zero probability of climate impacts so severe as to require immediate and rapid phase-out of anthropogenic greenhouse gas emissions. In Figure 5b, we examine the sensitivity of the most robust strategy to these assumptions. To make this figure, we find the most robust strategy as a function of: i) our expectations that immediate and draconian emissions reductions (we call this the Drastic Reductions strategy) are the best response to the threat of climate change and: ii) the probability that the damages due to climate variability are high or increasing. Our finding that the most robust adaptive-decision strategy has moderate near-term emissions and is more sensitive to observations of abatement costs than to observations of damages, is true in the shaded region on the left of the figure where the likelihood of drastic future ranges from about 0% to 12%, depending on the likelihood of low climate variability.

4.2. Promoting Innovation

If climate change is indeed a serious problem, society will have to make significant reductions in greenhouse-gas emissions, on the order of 80% below extrapolations of current trends, by the end of the 21st century. Technology innovation will likely play a major role in any changes of this scale. A number of modeling studies have made this point simply by treating innovation as an exogenous influence on key parameters, such as the rate of autonomous energy efficiency improvements, the cost of low-carbon emitting technologies, or the rate of technology spillovers from developed to developing countries (see for instance, Dowlatabadi, in press and Edmonds and Wise, 1998). By considering a range of assumptions about these parameters, these studies confirm that over the long-term, sustained rates of either fast or slow

innovation can make virtually any emissions-reduction target either inexpensive or impossibly costly to meet.

This fact, and our argument that the most robust, adaptive-decision strategies may emphasize technology innovation, raises a key question -- what should policy-makers do to encourage such innovation? In practice, governments pursue a wide range of technology policies designed to improve the technology options for emissions reductions in the future. Such policies include supporting research and development, training individuals, funding demonstration projects, building infrastructure, disseminating information, and implementing a variety of tax credits and subsidies to encourage the use of new technologies. These programs often appear attractive both on substantive and political grounds, but their record in practice is mixed. In addition, economic theory suggests that in the absence of market failures, such policies are inefficient compared to policies designed to "get the prices right," such as carbon taxes or tradable permits. While there is wide agreement that governments ought to fund research and development (though less agreement on the extent to which this funding ought to focus on specific areas such as emissions-reducing technologies), it is not clear the extent to which climate-change policy ought to focus on getting the price of carbon right, developing technology policies designed to improve future options for emissions reductions, or some combination of both.

Our judgements about this question will depend on our expectations about the ability of policy to change the dynamics of technology diffusion. Many recent studies have begun to address these issues (Grubler, Nakicenovic, and Victor, 1999; Azar and Dowlatabadi, 1999), focusing on both the effects of learning-by-doing on technology diffusion (Grubler and Gritsevskiy, 1997; Mattsson and Wene, 1997; Anderson and Bird, 1992) and on the accumulation of knowledge that can lead to new innovation (Goulder and Schneider, forthcoming; Schneider and Goulder, 1997; Goulder and Mathai, 1997). In addition, much empirical work (e.g., Newell et al., forthcoming; Grubler, Nakicenovic, and Victor, 1999) is becoming available to inform these modeling efforts. The modeling efforts to date, however, have

largely focused on creating scenarios of the future and showing their sensitivity to a variety of assumptions about technology innovation. They have not used this information to adjudicate among alternative policies and, in particular, examine the desirability of technology policies. In our recent work, we have examined the conditions under which technology incentives should be a key building block of a robust, adaptive-decision approach to climate change. This is an interesting question in its own right, but it also sheds important light on the design of adaptive design strategies and on the types of information about technology futures that, in the absence of an ability to forecast, should prove most useful to policy analysis and policy-makers.

We have addressed these questions using our exploratory modeling methods and a model of technology diffusion which focuses on the social and economic factors that influence how economic actors choose to adopt, or not to adopt, new emissions-reducing technologies. As shown in Figure 6, we use an agent-based model of technology diffusion coupled to a simple macro model of economic growth. Used in this fashion an agent-based model is merely a stochastic mathematical function representing the dynamics of factors in the macro-model, such as energy intensity. Like any time series model, its parameters can be calibrated to reproduce real-world data. Such agent-based models are particularly useful, however, because they conveniently represent key factors influencing technology diffusion, and thus, policy choices, such as the heterogeneity of economic actors and the flows of imperfect information that influence their decisions. Each agent in our model represents a producer of a composite good that is aggregated as total GDP, using energy as one key input. Each time period the agents first choose among several energy-generation technologies and second, given their chosen technology, choose how much energy to consume. (That is, agents choose a production function and where to operate on that production function.)

We assume that agents pick a technology in order to maximize their utility, which depends on each agent's expectations about the cost and performance of each technology. The agents have imperfect information about the current performance of new technologies, but can improve their information based on their own past experience, if any, with the technology

and by querying other agents who have used it (Ellison and Fudenberg, 1995). The agents are also uncertain about the future costs of new technologies, which may or may not decline significantly due to increasing returns to scale. Agents estimate these future costs based on observations of the past rates of adoption and cost declines. Thus, the diffusion rate can depend reflexively on itself, since each user generates new information that can influence the adoption decisions of other potential users.

We write each agent's utility function using a risk-adjusted, Cobb-Douglas form

$$\langle U_{i,g,j}(t) \rangle = \langle \text{Perf}_{i,g,j}(t)^{\alpha_i} \rangle \langle \text{Cost}_{i,g,j}(t)^{1-\alpha_i} \rangle - \lambda_i [\text{Var}_{\text{Perf}}(t) + \text{Var}_{\text{Cost}}(t)] \quad (3)$$

where the first term, $\langle \text{Perf}_{i,g,j}(t)^{\alpha_i} \rangle$, is the expectation of agent i , in region g , at time t of the performance it will get from technology j . This term depends on information flows among the agents, which we model crudely as a random sampling process, where each agent queries some number of the other agents each time period. The second term, $\langle \text{Cost}_{i,g,j}(t)^{1-\alpha_i} \rangle$, is the expected cost of using the technology over its lifetime, derived from observations of past trends in usage and cost of the technology. This term depends on the potential for cost reductions due to increasing returns to scale. The third term represents the agent's risk aversion, taken as a function of the variance of the estimates of technology performance and future costs.

Equation (3) focuses on the heterogeneity of the agents' preferences, which is important in creating the possibility for early adopters and niche markets that can strongly influence the course of technology diffusion. First, different agents obtain different levels of performance from the same technology because technologies differ according to characteristics such as size of the equipment and ease of maintenance that matter to some users more than others (Davies, 1979). We represent this in our model with a distribution of performance factors for each technology across the population of agents. Economic actors also have different cost/performance preferences for new technologies. Potential early adopters are generally much more sensitive to

performance than price while late-adopters are often more price sensitive. We represent this in our model by allowing agents to have different values for the exponents α_i , representing the cost/performance tradeoffs, and in the risk aversion coefficient λ_i . Finally, in a world of imperfect information, different economic actors will have different expectations about the performance and cost of each technology. At the beginning of each of our simulations, each agent is assigned an expectation about the performance of each technology, which it can update as time goes by. Each agent's expectations of performance are private; that is, they apply only to that agent, since in fact each agent will gain a unique performance from each technology. The cost forecasts are public, that is, shared in common by all the agents, but each agent in general will have different planning horizons, determined by the remaining lifetime of the technology they are already using. Each type of heterogeneity considered in our model can significantly affect market shares and the dynamics of the diffusion process (Bassanini and Dosi, 1998).

This model requires data on the social and economic context of technology diffusion, quite different from that generally demanded by other climate-change-policy models and often supplied by technology forecasts. Each of the technologies in our study is represented by three factors: the cost (which can drop over time due to increasing returns), the carbon intensity (the quantity of CO₂ emitted when generating one unit of energy), and the performance. The cost and the carbon intensity are intrinsic to the technology, and are treated in our model similarly to other models. However, the performance, as noted above, depends on the agents using it. In addition, we need information on the different preferences and current expectations of the users and potential users of the new technologies. Thus, we are interested in data which show how the performance of a technology varies across many different types of users and, in particular, how different key segments of the market along the technology adoption life cycle (Moore, 1995) -- the early adopters, early and late majority, and laggards -- may perceive, judge, and use the technology. In our model we treat these factors crudely, assuming in each case that preferences, expectations, and performance are distributed normally across our population of agents. Using these simple assumptions we can draw

what seem to be important, but general, policy conclusions. More detailed policy recommendations would probably require better information about such factors.

In addition to considering the agents' individual technology and energy consumption choices, we also need to consider the aggregate impact of the agents' actions. We capture these effects with a very simple representation of economic growth. Assuming a world economy in a *steady state*, we write the Gross Domestic Product (GDP) in each of two world regions, the OECD and the Rest of the World (ROW), with the difference equation, where output per capita grows at some exogenous rate that can be modified by changes in the price of energy and any damages due to climate change. Thus,

$$\text{GDP}_g(t) = \text{GDP}_g(t-1) \left[1 + \gamma_g - \phi_{xg} \cdot \Delta C_{\text{eng}}(t) - \phi_{sg} \cdot C_{\text{sub}}(t) \right] \left[1 - \text{impacts}(t) \right], \quad (4)$$

where γ_g represents the exogenous growth rate in the regions $g = \text{OECD}$ and ROW ; ϕ_{xg} and ϕ_{sg} are the elasticities in energy prices, including the costs of energy-producing technologies, any carbon tax imposed in order to reduce CO_2 emissions, and the cost of the subsidy to changes in economic growth; and impacts due to climate change are given by a simple polynomial function of the atmospheric concentrations of greenhouse gases, $\text{impact}(t) = \kappa_0 [\text{Conc}(t)/\text{Conc}(1765)]^{k_1}$, using a simple difference equation to relate concentrations to emissions (Cline, 1992 and Nordhaus and Yang, 1996). The aggregate decisions of the population of agents affect the cost of energy, the emissions, and thus, the climate impacts. In turn, the rate of GDP growth affects the number of new agents with new expectations and no sunk capital cost and, thus, the decisions of each individual agent.

Having defined our model, we can use it to create the landscape of plausible futures by defining ranges for the model inputs and constraints on the model outputs. On the micro (bottom-up) level, we confine ourselves to parameters describing only three generic types of technologies: high-emission-intensity systems, such as coal-fired power plants, which at present provide the bulk of the world's energy; medium-emissions-intensity systems,

such as natural-gas-powered combustion, which provide a significant minority of the world's energy; and low-emissions-intensity systems, such as renewable, biomass, and/or new nuclear power facilities, which at present are not in widespread use, but may be significant energy sources in the future. This is clearly a very coarse grouping, similar to that used by Grubler, Nakicenovic, and Victor (1999), but it is sufficient to draw policy conclusions about the appropriate mix of carbon taxes and technology subsidies. We choose a range of cost and performance parameters for these three generic energy technologies as well as parameters describing the behavior of the population of consumers of these energy technologies based on averages of data describing technology systems currently with significant market penetration and standard forecasts of technologies with potential significant future market penetration (Tester et. al, 1991; Manne and Richels, 1992).

On the macro (top-down) level, we use data for a variety of parameters describing the growth of the economy and its response to changes in the cost of energy and damages due to climate change (World Bank, 1996; Dean and Hoeller, 1992). These parameters include the exogenous rate of economic growth, the elasticity of economic growth with respect to the cost of energy, the elasticity of economic growth with respect to the cost of the subsidy, the parameters characterizing damages and the concentration of carbon in the atmosphere, the parameters defining the energy-demand functions, and those used to simulate exogenous improvements in energy efficiency.

We can now search across the thirty-dimensional space of model inputs that plausibly represent the micro-level data, looking for those combinations that give plausible, macroscopic model outputs. Given the simplicity of the model and the types of data readily available, we chose three constraints on model outputs: current (1995) market shares for energy technologies; current levels of carbon emissions, energy intensities, and carbon intensities; and diffusion rates no faster than 20 years (from 1% to 50% penetration). The first two constraints guarantee that our model is consistent with current data. The third, limiting diffusion to a rate no faster than the rates historically observed for energy technologies (Grubler, 1990), forces the

model to be consistent with one of the historically observed patterns of technology diffusion.

We next generate the most expansive ensemble we can produce of model input parameters consistent with these constraints. Using a genetic search algorithm, we generated 1,611 such sets of input parameters, covering a wide variety of assumptions about key parameters such as the level of increasing returns to scale, agents' heterogeneity – represented by the distribution of values for the coefficients in Eq. (3) across the population of agents, uncertainty regarding new technologies, and future damages due to climate change. This set does not include points with very small levels of uncertainty and heterogeneity regarding expectations about the performance of new technologies (such points do not satisfy the constraint on initial market shares), sets of points where the agents' utility functions are largely insensitive to costs, and sets of points with both very high learning rates and levels of increasing returns to scale.

We can now ask the question as to whether an adaptive-decision strategy ought to use technology incentives and carbon taxes, or carbon taxes alone, in order to address the threat of climate change. We posit two strategies, "Taxes Only" and "Combined Strategy", whose performance we compare across the landscape of plausible futures. (We also considered no response and a strategy using only technology subsidies. Neither of these very attractive compared to the other two.)

As the name implies, the "Taxes Only" strategy employs only a carbon tax, whose level can change over time in response to observations of the rate of economic growth and the damages due to climate change, as shown in Figure 7. The tax begins at some initial level per ton of emitted carbon and grows at a fixed annual rate. However, if in any year the cost of the tax is greater than the marginal cost of emissions of carbon dioxide (expressed as a percentage of GDP), the tax remains constant. Similarly, if in any year the global economy growth rate is below some minimum rate, the tax returns to its initial level, from which it can begin to grow again. This description of a steadily growing tax is consistent with the optimum taxes described in the

literature (Goulder and Schneider, 1998), and the stopping condition represents a way in which political conditions may force a tax to terminate.

The "Combined Strategy" uses both the carbon tax and a technology subsidy, which can change over time in response to observations of the market share of low-emitting technologies, as shown in Figure 7. The subsidy begins at some initial level, expressed as percent of the cost of the subsidized technology. This subsidy stays at a constant level over time until either the market share for low-emitting technologies goes above a threshold value or the market share fails to reach a minimum level after a certain number of years. If either of these conditions is met, the subsidy is permanently terminated. Thus, we assume the subsidy is terminated once policy-makers observe that the technology succeeds or that it fails to diffuse by some deadline.

These tax and incentive policies are together characterized by a total of seven parameters -- the beginning levels of the tax and subsidy, the annual increase in the tax, the minimum level of economic growth needed to maintain the tax, the maximum market share which terminates the subsidy, and the minimum market share over what time period that terminates the subsidy. We choose the particular value of these parameters used to define our "Taxes Only" and our "Combined Strategy" by searching for the best tax and the best subsidy strategy at the point in uncertainty space characterized by the average value for each of the model input parameters. The tax starts at an initial value of \$100/ton carbon and grows at 5% per year. The initial subsidy is 40% of the cost of the low-emitting technology. The subsidy terminates if the low-emitting technology reaches 50% market share, or fails to reach 20% market share after 15 years. This process is a crude approximation to the procedure used in LSBA, in which we found the best strategy for many different states of the world. We are currently applying the LSBA procedure to a comparison of tradable permits and technology incentives. The general conclusions presented here seem to hold.

In order to compare the performance of the alternative strategies, we calculate their performance for each of a large number of different states of the world, looking for those that distinguish one policy choice from another.

There are too many dimensions of uncertainty (30) in this model for an exhaustive search, so we used knowledge about a system and the goals of the analysis to summarize a very large space with a small number of key scenarios. We performed a Monte Carlo sample over the thirty-dimensional space, used econometrics (Aoki, 1995) to look for those input parameters most strongly correlated with a key model output of interest (GHG emissions after fifty years), and used variations across these key parameters to define scenarios. (This is a version of critical factors analysis, as in Hope (1993) and in Kalagnanam, Kandlikar, and Linville, 1998.) After making hypotheses about policy recommendations based on this reduced set of scenarios, we then tested our results by launching a genetic search algorithm (Miller, 1998) across the previously unexamined dimensions looking for counter examples to our conclusions.

Figure 8, a typical result of such comparisons, shows the regret of the Tax-Only strategy (dashed line) with the regret of the Combined strategy (solid line), as a function of the heterogeneity of the agent population. For this figure we have assumed moderately increasing returns to scale, a moderate level of social interactions, and moderate damages due to climate change. The figure shows that the Tax-Only strategy is preferable in a world where the agents are homogeneous. As the heterogeneity of the agents' preferences increases, the Combined tax and subsidy strategy quickly become more attractive.

The Combined Strategy becomes less costly than Taxes-Only due to the competition between two effects. Carbon taxes reduce emissions by slowing economic growth and inducing agents to switch to low-emitting technologies. Subsidies slow growth slightly, and may or may not induce agents to switch to low-emitting technologies. When the agent population has low heterogeneity, there are few potential early adopters and the tax is more efficient than the subsidy. Increasing heterogeneity favors the Combined Strategy, because it creates a number of potential early adopters that the subsidy will encourage to use the new, low-emitting technology which, in turn, generates learning and cost reductions that benefit society as a whole.

But the Subsidy-Only strategy never dominates in our simulations, because when the heterogeneity is large, there are always a substantial number of late-adopters who will reduce emissions only when faced with a tax.

We have considered a large number of results such as those in Figure 8, which we summarize in Figure 9. The figure shows the expectations about the future that should cause a decision-maker to prefer the Tax-Only strategy to the Combined Strategy. The horizontal axis represents the range of expectations a decision-maker might have for how likely it is -- from very unlikely on the left to very likely on the right -- that factors such as the potential number of early adopters and the amount of increasing returns to scale will significantly influence the diffusion of new technologies.⁴ The vertical axis represents the range of expectations a decision-maker might have that there will be significant impacts due to climate change (greater than 0.3% of the global economic product). The figure shows that the Combined Strategy dominates even if decision-makers have only modest expectations that impacts from climate change will be significant and that information exchange and heterogeneity among economic actors will be important to the diffusion of new, emissions-reducing technologies.

Our results are consistent with those of other models of induced technology change introduced in the climate-change literature in recent years. Gritsevski and Nakicenovic (1999) use a stochastic optimization technique to examine the effects of induced technology learning, uncertainty, and increasing returns, and technology clusters, but without heterogeneous preferences. They find a wide, bimodal range of “basecase” emissions scenarios and argue for near-term investment in emissions-reducing technologies focused on clusters of such technologies. Goulder and Schneider examine greenhouse-gas-abatement policies using a general equilibrium model for the United States that takes into account incentives to invest in research and development, knowledge spillovers, and the functioning of research and development markets. They find that the tax should be accompanied by a subsidy *only* when there are spillover benefits from research and development. The work presented here shows that this result can be more general. Many types of spillovers, such as those resulting from

increasing returns to scale, network externalities, or *non-R&D knowledge spillovers* from users to non-users of new technologies may justify a subsidy. However, we find that the level of spillovers that justify the subsidy depends on the degree of heterogeneity of agents preferences and their attitude towards risk, and that, all other things being equal, the critical level to justify the subsidy decreases as heterogeneity and risk aversion increase.

5. Conclusions

When viewed as an adaptive-decision strategy, one that can evolve over time in response to observations of the climate and economic systems, climate-change policy has important implications for new energy technologies and the role of forecasts of these technologies. In the work described here, we examined several facets of this broad question. First, we argue that an adaptive-decision approach to climate change will be significantly more robust in the face of deep uncertainty than an approach that does not adapt. We found that the near-term steps in such a robust, adaptive-decision response to climate change should likely include actions explicitly designed to encourage the early adoption of new, emissions-reducing technologies. Such actions are justified if there are significant social benefits of early adoption beyond those gained by the adopters themselves, as is the case when there is heterogeneity in economic actors' cost/performance preferences for new technologies, when new technologies have increasing returns to scale, or when potential adopters can learn from others about the uncertain performance of new technologies. These results are consistent with most other climate-policy analyses which include both significant uncertainty and technological change that responds to policy choices and economic signals.

In addition, we examined the impacts of climate variability on near-term policy choices and found that the future course of emissions-reducing technologies may be one of the key indicators that should shape the evolution of an adaptive-decision strategy. The ability of future innovation to radically reduce emissions of greenhouse gases during the 21st century is one of the key uncertainties facing climate-change policy, perhaps more so than our uncertainty in the response of the climate system to those emissions.

Particularly in the presence of climate variability, over time, policy-makers may gain more definitive information about the potential of emissions-reducing technology than they will about the damages due to anthropogenic influences on the climate. Over the next decades then, the most reliable information policy-makers can use to adjust emissions constraints on greenhouse gases and other policy actions may come from indicators of society's ability to reduce such emissions at low cost.

Combining these findings supports the broad view that in the near-term, climate-policy ought to be more focused on improving society's long-term ability to reduce greenhouse-gas emissions than on any particular level of emissions reductions. As a long-term goal, the Framework Convention on Climate Change aims to stabilize atmospheric concentrations of greenhouse gases at some safe, currently unknown and unspecified, level. As a first step, the Kyoto Protocol requires the developed countries to reduce their greenhouse-gas emissions to specific targets by the end of the next decade. Both advocates and opponents usually see these emissions caps as hedging actions, designed to decrease the severity of any future climate change by reducing the concentration of greenhouse gases in the atmosphere. Thus, much of the climate-change debate, particularly in the United States, takes the particular level of these goals very seriously, and revolves around whether or not such targets are justified, and seeks to determine the least expensive means to meet them.

An adaptive-decision approach, combined with the uncertain, but large potential of the various technologies described in this volume, suggests that the Kyoto targets are in fact shaping actions, whose most significant impact may be making future emissions reductions, if necessary, less costly and easier to implement, rather than reducing atmospheric greenhouse-gas concentrations. Goals and forecasts often serve this type of shaping function. For instance, the leadership of organizations, such as private firms, often set goals for sales or profits intended to motivate their employees. The leaders would be disappointed if the goals were consistently met, for that would be an indication they had not asked their employees to strive hard enough. Conversely, the firm's financial planners will likely produce conservative

sales and profit forecasts and be disappointed if reality did not exceed their expectations. Similarly, the Kyoto targets may be most important as a signal to numerous actors throughout society that they should take the climate problem seriously and plan their actions and investments accordingly.

To some extent, Kyoto has already been effective in this regard. For instance, many firms are beginning to reduce their own emissions, factor potential climate change into their internal planning, and invest in research that would help them reduce emissions more in the future. Traders have begun to implement markets for trading carbon-emissions rights. Governments have begun to construct the institutions necessary to implement international regulation of greenhouse-gas emissions. It is unclear and, for an adaptive-decision strategy, perhaps irrelevant, whether or not these actions will result in Kyoto emissions-reductions targets being met. Rather, the question becomes what are the means to induce technological change that best balance between the possibility that such innovations will soon be needed to address severe climate change and the possibilities that they will not be needed or that the technologies will not meet their promise.

In this context, technology forecasts, to be most useful for informing the design of adaptive-decision strategies, should provide a broader range and different types of information than is often the case. Typically, an energy-technology forecast describes the future costs of the power produced by the system, estimates the maximum amount of society's energy needs the technology might satisfy, and reviews the technology's pros and cons. Often these predictions stem from an analysis of fundamental physical or engineering constraints and/or from analogies from similar systems. For instance, forecasts of the potential generation capability of future photovoltaic energy systems might be based on comparisons of the current and theoretical efficiencies of various types of cells and a survey of the solar insolation and land available for photovoltaic installations. Forecasts of the future costs of these systems might compare past rates of cost reductions to limits imposed by the basic costs of materials (that is, setting the design, production, and installation costs to zero). Sometimes these technology forecasts are given as best estimates; sometimes as a range of scenarios.

Generally they are conceived as an answer to the question: how important to society will this technology be in the future?

Such information is useful to policy analyses of climate change. But it is also contrary to much of what we know about processes of technology change, especially over the time-scales relevant to the climate problem. In the long term, we know costs and market shares of particular technologies are impossible to predict with any accuracy. We also know that social and economic factors, generally under-emphasized in technology forecasts, are among the primary determinants of the fate of new technologies. The unpredictability of such social and economic factors is a reason that technological change itself is impossible to predict. As one example, a recent study found that cost projections for renewable-energy systems had been reasonably accurate over the last twenty years, while predictions of market share have proved significantly too optimistic (RFF ref). The reasons are not entirely clear, but cost reductions may have been helped by rapid advances in materials and information technologies occurring outside of renewable-energy development, and the market shares have been hindered by the fossil-energy prices being much lower than past forecasters imagined possible.

Nonetheless, these social and economic factors impose common patterns on the processes of technology diffusion and can play an important role in constraining future scenarios, thus helping us choose the best near-term steps in an adaptive-decision strategy. For instance, we know that new technologies generally experience lifecycles that begin with periods of rapid improvement in cost and performance, continue with a period of mature incremental improvements, and end with a period of replacement by newer technologies. A technology's early stages are characterized by many firms experimenting with alternative designs, followed by a period of consolidation over a single design and competition of better means to produce and distribute that design (Utterback, 1994). We know there are common patterns of entry and exit as technologies compete against one another in the market place, although radical new technologies sometimes replace established competitors, the old technologies can sometimes crush the new

technologies with a burst of incremental improvements. We know there are network effects in which improvements, or the lack thereof, in one technology area can impact improvements in other areas.

Our analysis gives one example of how such patterns, and information about the social and economic factors that influence them, can be used to help design adaptive-decision strategies. We find that expectations about factors such as the heterogeneity of the potential adopters of a new technology – the availability of niche markets – and the speed with which information about experience with a new technology travels through a society – the networks and industry structure among these niche markets – are important in determining the best policy choices. In our analysis we used only crude estimates for such processes. Technology forecasts that provide information on factors such as the size of niche markets, the types of networks in which these potential early adopters reside, and the types of cost/performance preferences they display, could greatly enhance analysis of this type. Forecasts that examine how these factors might play out in a variety of diffusion scenarios would be even more useful. Such information might be seen as less rigorous than the hard, physical constraints on predicted cost and total deployment generally offered in technology forecasts. However, good descriptions of early markets and potential diffusion paths are a more realistic goal than predicting the long-term course of new technologies and, as described in this chapter, provide a more solid basis on which to design responses to climate change that adapt over time and, thus, are robust across a wide range of plausible futures.

REFERENCES

Anderson, D. and C.D. Bird, 1992: “Carbon Accumulation and Technical Progress – a simulation case study of costs,” Oxford Bulletin of Economics and Statistics, 54: 1-29.

Aoki, M. 1986: American Economic Review.

Aoki, M., 1990: Journal of Economic Literature.

Azar, C. and Hadi Dowlatabadi, 1999: "A Review of Technical Change in Assessments of Climate Policy, Annual Review of Energy and Environment.

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

Bassanini, A. P. and G. Dosi, 1998: *Competing Technologies, International Diffusion and the Rate of Convergence to a Stable Market Structure*, Interim Report from the International Institute for Applied Systems Analysis, (1998).

Casan, E.A., M. G. Morgan, and H. Dowlatabadi, 1999: "Mixed Levels of Uncertainty in Complex Policy Models," Risk Analysis, 19, 1, 33-42.

Clark, W. C., 1980: "Witches, Floods, and Wonder Drugs: Historical Perspectives on Risk Management," in *Societal Risk Assessment: How Safe is Safe Enough?* eds. Richard Schwing and Walter Albers, Jr., Plenum Press, New York.

Cline, W., 1992: *The Economics of Global Warming*, Institute for International Economics, Washington, DC.

Davies, 1979: *The Diffusion of Process Innovations.*, Cambridge University Press, Cambridge, 1979.

A. Dean A., and P. Hoeller, 1992: *Costs of Reducing CO2 Emissions. Evidence From Six Global Model*, Working Papers, Department of Economics, no 122. OECD, Paris.

Dewar, James A, Carl H. Builder, William M. Hix, and Morlie H. Levin, 1993: Assumption Based Planning: A Planning Tool for Very Uncertain Times, RAND, Santa Monica, CA, MR-114-A.

Dowlatabadi, Hadi, 1999: "Adaptive Management of Climate Change Mitigation: a strategy for coping with uncertainty, paper presented at the Pew Center Workshop on the Economics and Integrated Assessment of climate change, July, Washington, D.C.

Dowlatabadi, Hadi, in press, "Sensitivity of Climate Change Mitigation Estimates to Assumptions about Technical Change," *Energy Economics*.

Edmonds, J. and M. Wise, "The value of advanced energy technologies in stabilizing atmospheric CO2," in *Cost-Benefit Analyses of Climate Change*, F.L. Toth (ed.), Birkhauser, 1998.

Ellison, G. and D. Fudenberg, 1995: *Quarterly Journal of Economics*, Vol. 440 93-125.

Goulder, L.H. and S. H. Schneider, 1998: "Induced Technological Change, Crowding Out, and the Attractiveness of CO₂ Emissions Abatement", *Resource and Environmental Economics* (submitted).

Goulder, L.H. and K. Mathai, 1997: *Optimal CO₂ abatement in the presence of induced technological change*, Working Paper. Stanford University, 1997.

Gritsevskiy, A. and N. Nakicenovic, 1999: "Modeling Uncertainty of Induced Technological Change," paper presented at the Pew Center Workshop on Economics and Integrated Assessment of Climate Change, July 21-22, Washington, DC.

Grubb, M., T. Chapuis and M. Ha-Duong, 1995: *Energy Policy*, Vol. 23 417-431.

Grubler, A., 1990: *The Rise and Fall of Infrastructures*, Physica Verlag, Heidelberg.

Grübler, A. and A. Gritsevskii , 1998:*Environmentally Compatible Energy Strategies Project*. IIASA, Laxenburg, Austria.

Grubler, A., N. Nakicenovic, and D.G. Victor, 1999: "Dynamics of energy technologies and global change," *Energy Policy*, 27, 247-280.

Grubler A. and A. Gritsevskiy, 1998: "A Model of Endogenized Technological Change Through Uncertain Returns on Learning," submitted to the *Economic Journal*.

Hammitt, J. K., Lempert, R. J., and Schlesinger, M. E., 1992: A sequential-decision strategy for abating climate change. *Nature*, 357, 315–318.

Hasselmann, K.: 1976, 'Stochastic climate models, Part 1: Theory', *Tellus* **28**, 473-485.

Hope, C., J. Anderson, and P. Wenman, 1993: "Policy Analysis of the greenhouse effect: An application of the PAGE model," *Energy Policy*, 21, 3.

Katz, R., and Brown, B. G.: 1992, "Extreme Events in a Changing Climate: Variability is More Important than Averages," *Climatic Change* 21: 289-302.

Lave, L. and Dowlatabadi, H.: 1993, 'Climate Change Policy: The Effects of Personal Beliefs and Scientific Uncertainty,' *Environ. Sci. Technol.* 27, 1692-1972.

Lempert and Schlesinger, *Robust Strategies for Abating Climate Change*, *Climatic Change*, forthcoming.

Lempert, R. J., M. E. Schlesinger, Steven C. Bankes, Natalia G. Andronova, 2000: "The Impact of Variability on Near-Term Climate-Change Policy Choices", *Climatic Change*, (forthcoming).

Lempert, R. J., M. E. Schlesinger, Steven C. Bankes, 1996: 'When We Don't Know the Costs or the Benefits: Adaptive Strategies for Abating Climate Change,' Climatic Change, 33, pp. 235-274.

Lindzen, R. S.: 1990, 'Some Coolness about Global Warming,' Bull. Amer. Meteorol. Soc., 71, 288-299.

Mattsson, N. and C.O. Wene, 1997: "Assessing new energy technologies using an energy system model with endogenised experience curves," Journal of Energy Research, 21: 385-93..

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

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

March, James G., 1994: *A Primer on Decision-Making*, Free Press, New York.

Marengo, L., 1992: *Journal of Evolutionary Economics*, Vol. 2. (1992).

Matheson, J.E. and R.A. Howard, 1968: An Introduction to Decision Analysis, SRI.

Mattsson, N., 1997: *Internalizing Technological Development in Energy Systems Models*, Thesis paper, Energy Systems Technology Division, Chalmers University of Technology. Goteborg, Sweden.

Mendelsohn, R., W. Nordhaus, and D. Shaw: 1994, "The Impact of Global Warming on Agriculture: A Ricardian Analysis" American Economic Review 84: 753-771.

Miller, John H. 1998: "Active Nonlinear Tests (ANTs) of Complex Simulations Models," *Management Science* 44:6:820--30

Morgan, M.G. and Dowlatabadi, H, 1996: "Learning from Integrated Assessments of Climate Change," *Climatic Change*, 34, 337-368.

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

Morita, Tsuneyuki, and Hae-Cheol Lee, "Appendix to Emissions Scenarios Database and Review of Scenarios," *Mitigation and Adaptation Strategies for Global Change*, vol. 3, pp. 121-131, (1998).

Munich Re (Muncherner Ruckversicherungs-Gesellschaft, D-80791 Munchen, Germany (<http://www.munichre.com>), 1997).

Nakicenovic, Nebojsa, Nadejda Victor, and Tsuneyuki Morita, 1998: "Emissions Scenarios Database and Review of Scenarios," *Mitigation and Adaptation Strategies for Global Change*, vol. 3, pp. 95-120.

Nakicenovic, N. and et al., 1999: *IPCC Special Report on Emissions Scenarios*. Cambridge University Press, Cambridge.

Newell, R. G., A. B. Jaffe, and R. N. Stavins, "The Induced Innovation Hypothesis and Energy-Saving Technological Change, *Quarterly Journal of Economics* (forthcoming).

Nordhaus, W. D. and Z. Yang, 1996: *The American Economic Review*, Vol.86 N4.

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

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

Park, George, and Robert J. Lempert, 1998: "The Class of 2014: Preserving Access to California Higher Education," RAND, MR-971-EDU.

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.

Repetto, R. and Duncan Austin, 1997: *The Costs of Climate Projection: A Guide for the Perplexed*, World Resources Institute.

Robalino, David A. and Robert J. Lempert, 1999: "Carrots and Sticks for New Technology: Crafting Greenhouse Gas Reductions Policies for a Heterogeneous and Uncertain World," *Integrated Assessment*, (forthcoming).

Rotmans, J. and H.J.M. de Vries, 1997: *Perspectives on Global Change: The TARGETS Approach*, Cambridge University Press, Cambridge.

Savage, L.J., 1954: *The Foundations of Statistics*, New York, Wiley.

Schlesinger, M. E., and N. G. Andronova, 2000: Maximum-likelihood estimation of climate sensitivity and anthropogenic sulfate forcing (in preparation).

Schlesinger, M. E., and X. Jiang, 1991: Revised projection of future greenhouse warming. *Nature*, 350, 219-221.

Schlesinger, M. E. and N. Ramankutty, 1992: Implications for Global Warming of Intercycle Solar-Irradiance Variations. *Nature*, 360, 330-333.

Schlesinger, M. E., X. Jiang and R. J. Charlson, 1992: Implication of anthropogenic atmospheric sulphate for the sensitivity of the climate system. In *Climate Change and Energy Policy: Proceedings of the International Conference on Global Climate Change: Its Mitigation Through Improved Production and Use of Energy*, Rosen, L. and R. Glasser (Editors). American Institute of Physics, New York, pp. 75-108.

Schlesinger, M. E. and N. Ramankutty, 1994: An Oscillation in the Global Climate System of Period 65-70 Years. *Nature*, 367, 723-726.

Simon, Herbert A., 1959: "Theories of Decision-Making in Economics and the Behavioral Sciences," *American Economic Review*, June.

Tester, J.W., D.O. Wood, and N.A. Ferrari, (eds), 1991: *Energy and the Environment in the 21st Century*, MIT Press.

Utterbeck, James M. *Mastering the Dynamics of Innovation*, Harvard Business School Press, 1994.

van Asselt, M. B. A. and Rotmans, J., 1996: "Uncertainty in Perspective," *Global Environmental Change*, 6(2), 121-157.

van der Heijden, K., 1996: Scenarios: The Art of Strategic Conversation, Wiley and Sons, 305 pp.

Watson, S. and Buede, D, 1987: Decision Synthesis, Cambridge University Press, Cambridge.

Wigley, T. M. L., R. Richels and J. A. Edmonds, 1996: Alternative emissions pathways for stabilizing CO₂ concentrations. *Nature*, 379, 240-243.

World Bank. *World Development Report*. Washington DC, 1996.

¹ There are, of course, additional debates over topics such as the extent to which the fast-growing developing countries ought to bare the burden of emissions reductions, the extent to which developed countries can pay others to make emissions reductions for them, and the best means to reward those firms that voluntarily commit to early actions.

² The N-year running average is given by $\Delta\bar{T}_N(t) = \frac{1}{N} \sum_{\tau=t-N+1}^t \Delta T(\tau)$.

³ Formally, we divide the states of the world into exclusive sets depending on whether DAL, ES, or drastic reductions (DR) is the better policy. (Drastic Reductions eliminates anthropogenic greenhouse gas emissions over the first half of the 21st century.) We assign a probability $(1 - p_{ES})(1 - p_{DR})$ to the states where DAL is better, $p_{ES}(1 - p_{DR})$ to the states where ES is better, and p_{DR} to the states where DR is better. The horizontal axis in Figure 5a spans $0\% \leq p_{ES} \leq 100\%$, with $p_{DR} = 0\%$. The static DAL and ES policies do not necessarily have zero expected regret at $p_{ES} = 0\%$ and 100% , respectively, because one of the adaptive-decision strategies may perform better across the DAL or ES states of the world.

⁴ This classicalness index $CI=\{0,i,\dots,4,5\}$ characterizes the triplet of parameters $(\beta_3, \vartheta, \sigma)_i$, $i \in CI$ where each parameter z of the triplet is given by: $z_i = z_{\min} + i \left(\frac{z_{\max} - z_{\min}}{5} \right)$, and $[z_{\min}, z_{\max}]$ determines the range of variation of the parameter.

Adaptive Strategy Dominates

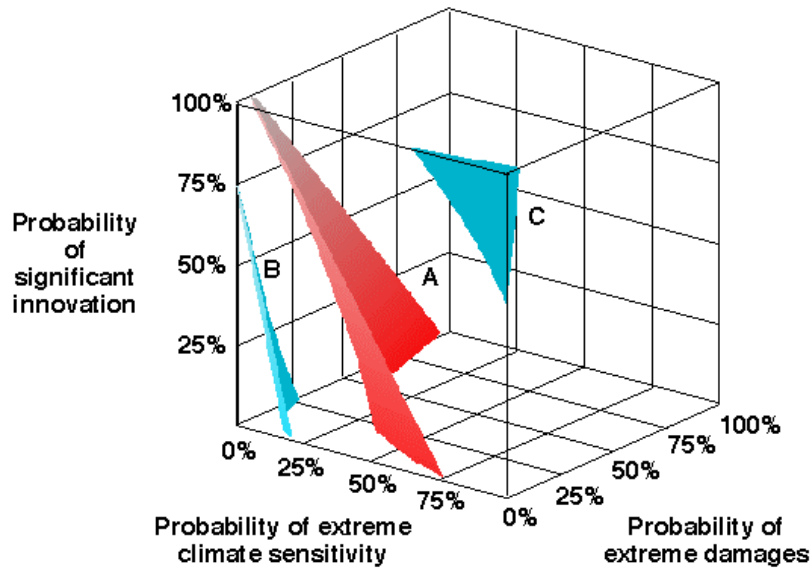


Figure 1: Adaptive decision strategies are the best approach to climate change policy.

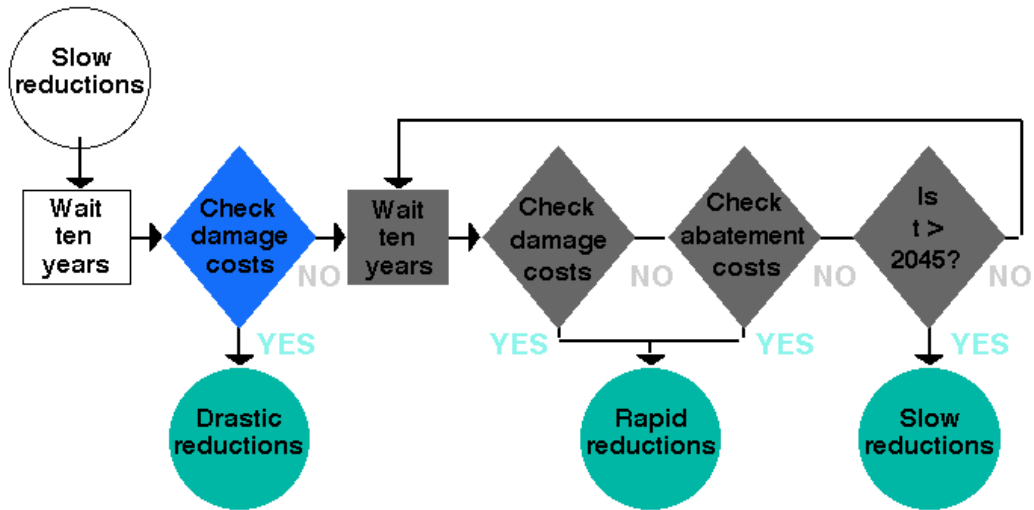


Figure 2: A simple adaptive-decision strategy

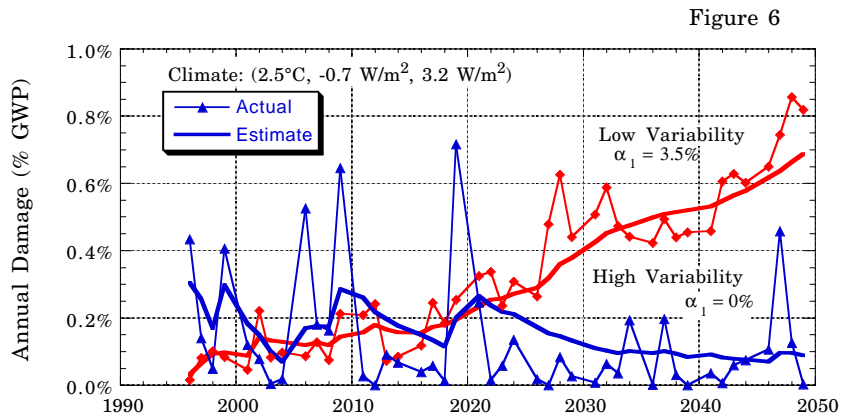


Figure 3: Model assumes decision-maker uses Bayesian estimator to extract any damage trend from the noise due to variability.

Figure 4

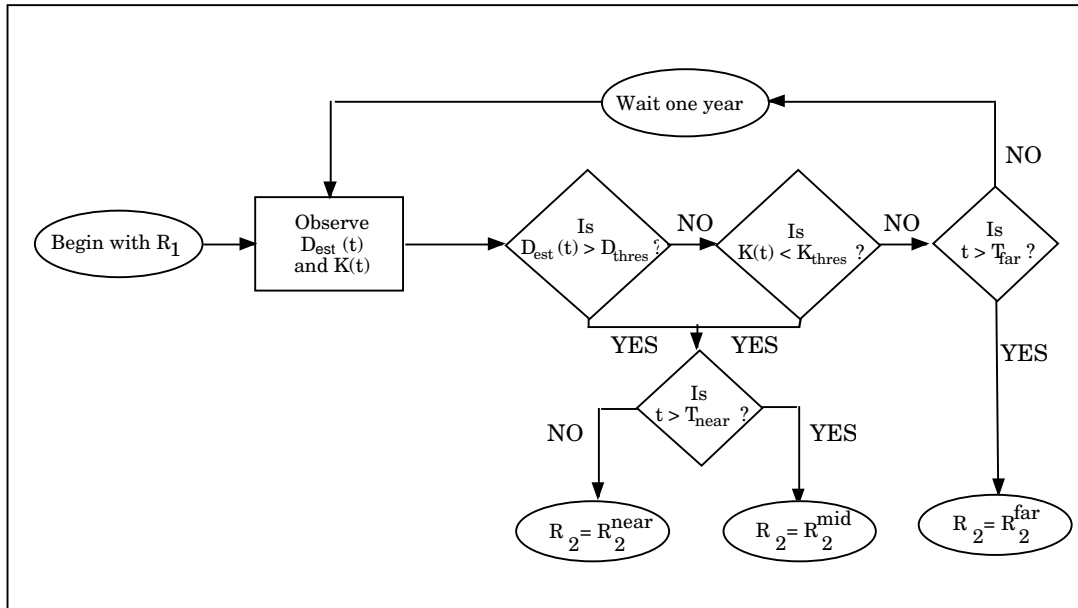


Figure 4: Adaptive Decision Strategy

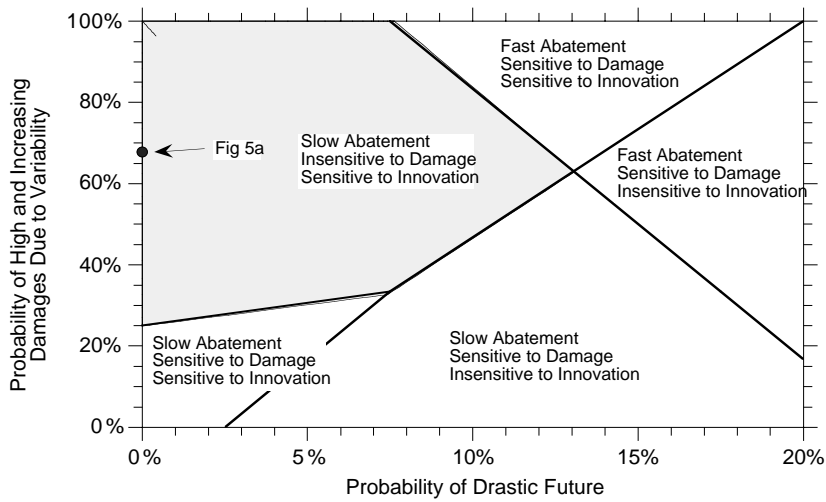
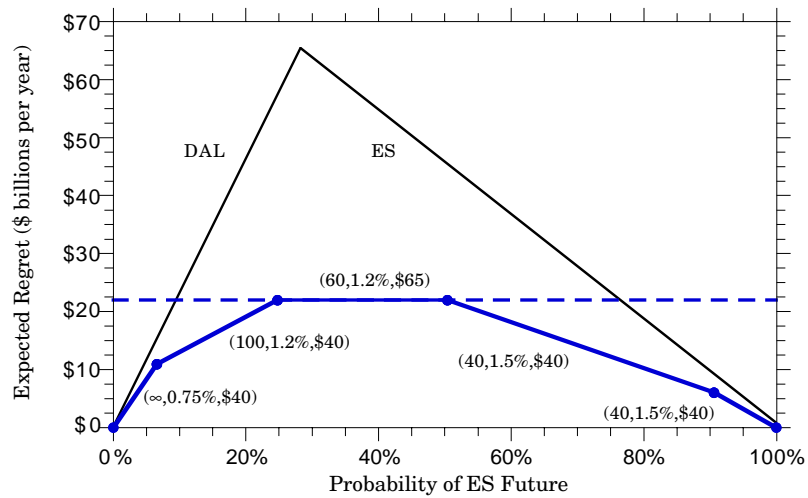


Figure 5: (a) Performance of alternative adaptive-decision strategies as a function of expectations about the future and (b) the most robust adaptive-decision strategy as a function of expectations about the probability of a Drastic Future and of high and increasing damages due to variability.

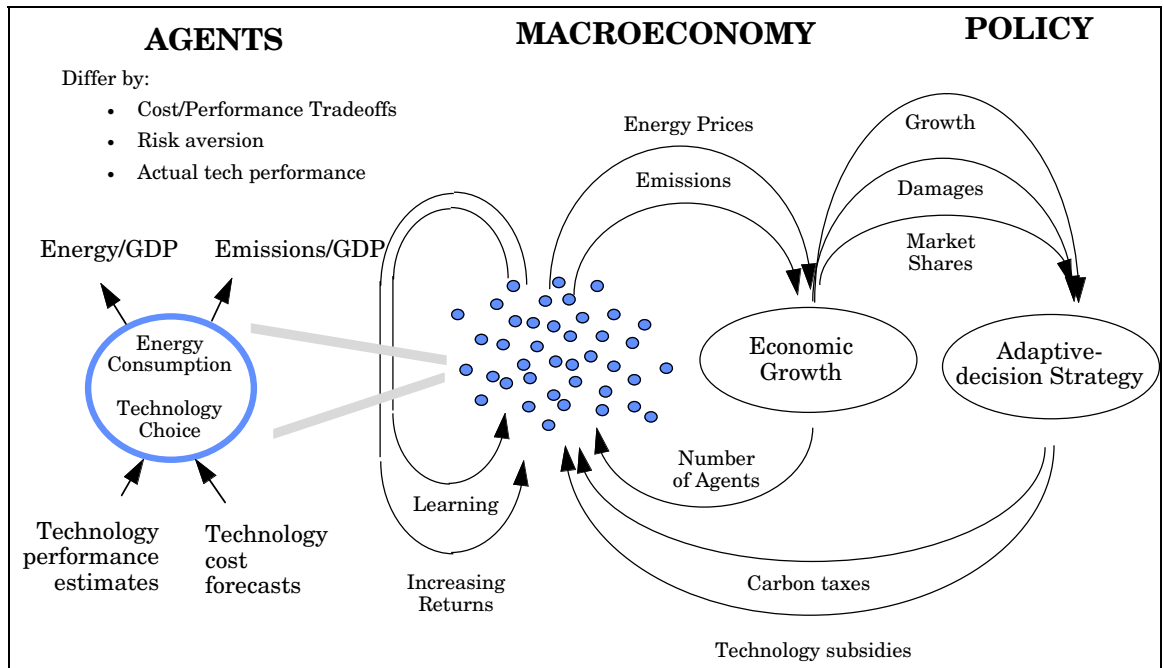


Figure 6: Agent-based model of technology diffusion used in this study.

Economic agents choose among alternative technologies on the basis of forecasts of cost and performance. The forecasts are influenced by learning among the agents and potential price decreases due to increasing returns to scale. The agents have heterogeneous initial expectations about technology performance and heterogeneous preferences for technology cost/performance tradeoffs. The agents' choices influence the level of energy prices and of greenhouse gas emissions, which both influence the rate of economic growth. Policy decisions about the level of carbon taxes and technology subsidies, which depend on observations of economic growth, damages and technology diffusion, also influence the agents' technology choices.

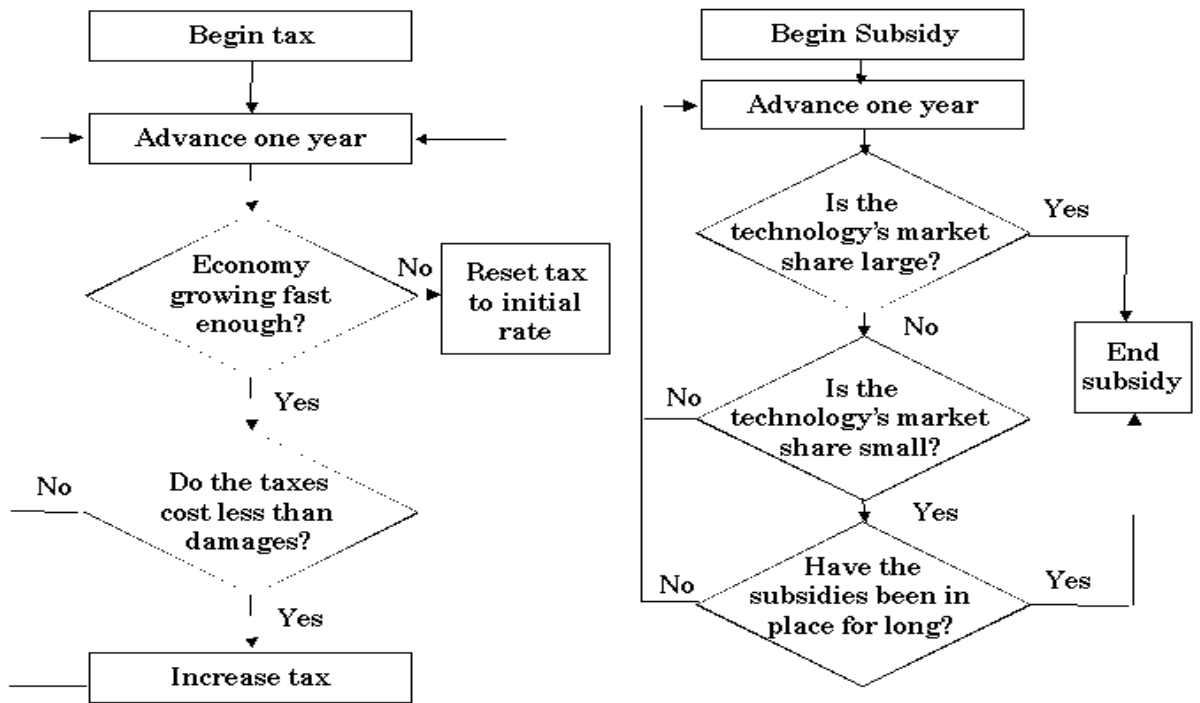


Figure 7: Adaptive-decision strategies for carbon taxes and technology subsidies

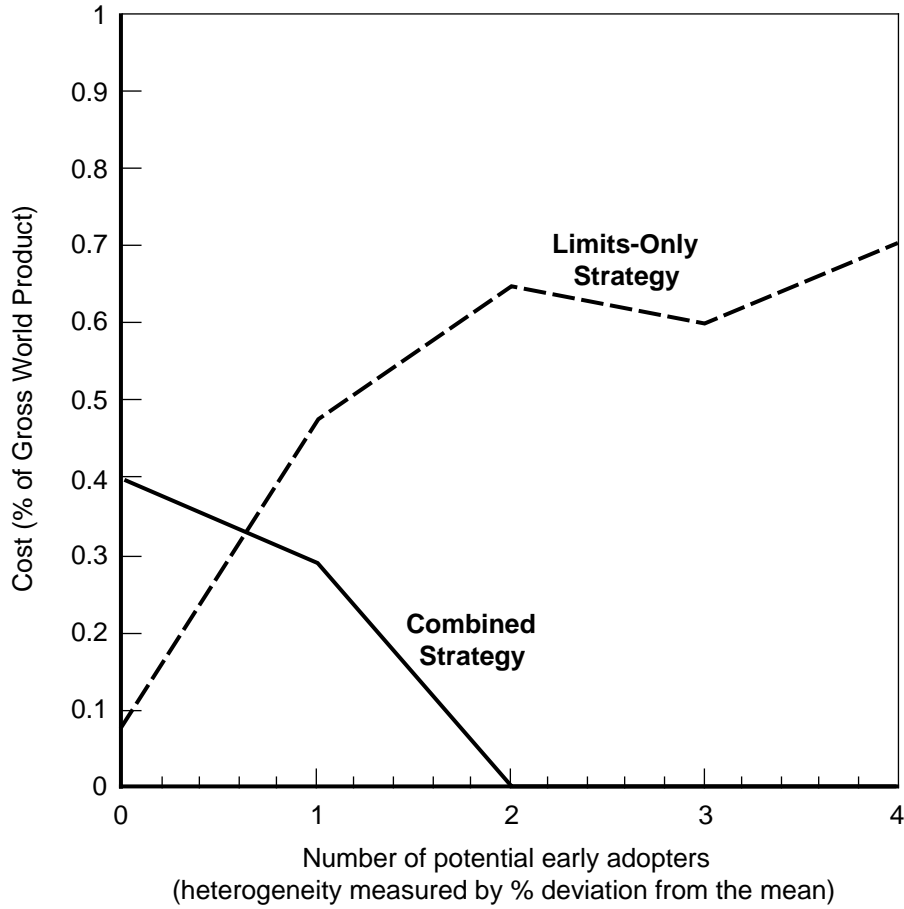


Figure 8: Expected regret of the Tax-Only (dashed line) and Combined tax and technology subsidy (solid line) adaptive-decision-strategies. The lines are displayed as a function of heterogeneity of the agents' population. All other input parameters are held constant at their mean values.

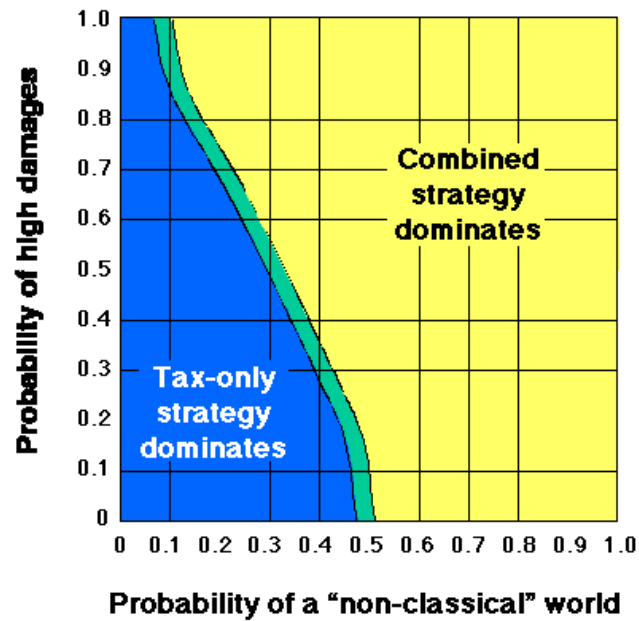


Figure 9: Regions in probability space where the expected GDP resulting from the Tax-Only strategy is greater than that of the Combined strategy as a function of the probability of a classical world and low damages due to climate change, as defined in the text.