ROBUST STRATEGIES FOR ABATING CLIMATE CHANGE

An Editorial Essay

People are of two minds about uncertainty. On the one hand, we crave predictions about the future. Virtually every newspaper, magazine, and trade publication has articles predicting some aspect of the future, and many people make good livings foretelling what the future will bring. The quest for prediction probably fills some deep human need. Even though the accuracy of most predictions has proven to be poor (Sherden, 1998), the future becomes less scary once given a name and a shape.

On the other hand, we usually recognize that predictions are inherently unreliable and have been extraordinarily creative in managing our affairs without them. From the folkways of farmers, to the basic tenets of most religions, to the rules of thumb that we use to navigate late 20th-century life, few provide predictions and most are designed to help the individual achieve some measure of success and fulfillment no matter what the future brings. The currently triumphant forms of social organization, democracy and the free market, are both explicitly designed to be self-correcting of errors and robust against unforeseen circumstances. Both assume people cannot predict the future and will make mistakes. Both probably owe much of their success to the truth of this view.

In recent years, researchers examining alternative policies to address the threat of climate change have become increasingly concerned about uncertainty. This is clearly appropriate, for few policy problems are dependent on such significant unknowns. Ascertaining the potential impacts of human activities on the immensely complex climate system and its related ecosystems is daunting enough. But climate-change policy must also concern itself with the state of society fifty and a hundred years hence. Will innovation drive the cost of non-emitting energy technologies below that of fossil fuels? Will new means of extraction and processing make coal the low-cost fuel of the future? Will society treasure the preservation of the present state of nature, or will our descendants¹ prefer the worlds that they themselves create? If we knew the answers to such questions, they would shape our choices about climate-change policy today.

The traditional framework for assessing alternative climate-change policies, which influences much climate-change-policy research and informs the thinking of many of the most sophisticated policy-makers, rests on the assumption that we can predict the future. In this framework, we begin with a set of alternative actions we might take; a model, often described mathematically, that allows us to describe the consequences of each action; and some metric, such as monetary units, that



Climatic Change 45: 387-401, 2000.

allows us to rank our relative preferences for various consequences. Analysts use this framework to predict the consequences of each action, and thus recommend the 'optimum' response, that is, the action that is better than all the alternatives. This framework treats uncertainty, when it is treated at all, by estimating the likelihood that different futures will come to pass. Rather than predict a specific outcome of any action, we summarize all the possible consequences of an action with an average (expected) value. The optimum policy becomes that alternative which, on average, performs better than all the others.²

This approach has numerous virtues. It has been used successfully for decades to solve a wide range of decision problems. It is supported by a deep body of theory and an elegant toolbox of mathematical techniques. But for many problems, such as those posed by climate change, such prediction-based policy analysis can be misleading, because its underlying premise of what we know about the future is not true. Perhaps more importantly, viable alternatives to this approach are now available. Rather than find the optimum strategy based on predictions, researchers using modern computer technology can now systematically and analytically evaluate alternative policies against a wide range of plausible scenarios and, thereby, directly address the real task that faces climate-change decision-makers – crafting strategies that are robust in the face of an unpredictable future.

1. The Problem with Prediction-Based Policy Analysis

Many research groups have recently attempted to forecast the cost to the United States of meeting its commitments under the Framework Convention on Climate Change, negotiated in Kyoto in December 1997. Some have argued that compliance with the treaty, which requires the U.S. to reduce its emissions of greenhouse gases by 7% below 1990 levels for the period 2008 to 2012, could cost up to 3% of GDP. Others have calculated that emissions reductions could improve GDP by a fraction of a percent (see, for instance, Repetto and Austin (1997), for a review of various estimates of the costs of greenhouse-gas abatement and the reasons why they differ). While some studies are better than others, these different projections largely reflect different, plausible assumptions about the future. In fact, we do not know what the Kyoto commitments will cost. The outcome depends on numerous factors – from the cost of various fossil fuels, the health of the economy, the progress of new technologies, to the efficiency with which government programs are put into place. We can predict none of these with any accuracy.

Estimating Kyoto-compliance costs is one of the easiest of the climate-change forecasts to make since the time horizon is only a little more than ten years hence. To predict the overall impacts of climate change requires not only forecasts of the behavior of the climate system itself, but also the evolution of ecosystems and social, technological, and economic systems decades into the future. People are often wildly wrong in predicting such social, technological, and economic change.

388

Technology forecasters, for instance, almost always miss the discontinuous technology changes, such as the advent of the passenger jet or the World Wide Web. Even the best market forecasters miss major shifts in consumer preferences. And the world's experts frequently fail to predict significant political and economic events, from the fall of the Soviet Union to the recent crash of Asian economies.

Any single scenario of the future is almost sure to be wrong. Prediction-based policy analysis attempts to preserve the idea of an optimum policy in the face of multiple, plausible scenarios by estimating the likelihood of each alternative future. Often this is done by specifying probability distributions for each of a number of input parameters to a model and propagating these probabilities through the model to estimate the probability of various future scenarios. For instance, researchers might estimate a 50% probability that the climate sensitivity is 2.5 °C and a 25% probability that it is 1.5 °C or 4.5 °C and, combined with estimates of the probability that future greenhouse-gas emissions will be high or low, calculate the likelihood of significant climate change over the next century. (The climate sensitivity is the equilibrium warming per unit of radiative forcing due, for example, to a doubling of carbon dioxide concentration in the atmosphere.) The optimum policy is then calculated based on these probabilities.

This approach can adequately address the problem of decision-making under uncertainty if the estimated probabilities are in some sense a 'correct' representation of the future. But often they are not. In order to estimate the probable outcomes of science that is not yet completed or future events that have not yet occurred, analysts often elicit the views of experts in relevant fields. For instance, policy researchers might ask climate scientists to estimate the likelihood that the climate sensitivity takes on certain values (Morgan and Keith, 1995). Or, they might consult the many technology experts who forecast the likelihood that coal or some set of non-greenhouse-gas-emitting alternatives will meet the energy needs of the 21st century (see, for example, Johansson et al. (1993) for a variety of projections for renewable energy sources). These elicitations are an excellent source of information. But there is real danger in regarding the resulting probabilities (called subjective probabilities) as a good foundation for defining optimum policies. People have many well-documented biases and failures in making such estimates. Moreover, there is often little information other than expert judgment on which to base such expectations. Some statisticians argue that subjective probabilities are always inappropriate because the concept of probability loses all meaning for future events that will happen only once. This is not our criticism. We acknowledge the alternative Bayesian view that meaningful subjective probabilities are possible, and in some cases, represents the best means to adjudicate among different policy options. But, we argue that using these probabilities to define optimum policies can lead to two critical mistakes in the assessment of climate-change policies.³

First, the concept of an optimum policy assumes a single, rational decisionmaker whose expectations about the future are well-approximated by a single set of probabilities. But society contains a multitude of actors, each with their own

expectations about the future. Thus, no optimum policy so based is likely to support the consensus needed for political action. Many different stakeholders are affected by the climate-change problem, and they hold very different views about the climate-change future. In part, such differences are due to different interpretations of the available scientific evidence. But, to a significant extent, they are also due to differing interests. Decision-makers and other political actors often understand that particular expectations support particular policies and that the available science supports a wide range of plausible futures. Contending stakeholders will do their best to choose divergent subjective probabilities that support the particular position they wish to hold on ideological, financial, or other grounds.⁴

Second, optimum policies can be brittle in the face of catastrophes, pleasant surprises, or other high-consequence, low-probability events. The decision-analysis literature often distinguishes among those cases (Morgan and Henrion, 1990; IPCC, 1996; NRC, 1997) in which we know the underlying system model, the loss function, and the subjective probability distributions of key, uncertain model parameters and those cases where the system model, loss function, and/or probability distributions are in doubt. In the former case, it can make sense to speak of an optimum response to the uncertainty. But in the latter, the case of deep uncertainty, such an optimal strategy may be misleading. For instance, many climate-change-policy studies advocate optimum hedging strategies with slight increases in near-term emissions reductions or slight increases in carbon taxes as a response to low-probability catastrophes. But such results do not inspire great confidence, because they often depend sensitively on assumptions about things we know very little about, such as the magnitude of potential catastrophes and their likelihood.

Real-world decision-makers often think about hedging quite differently from the way envisioned by the prediction-based policy-analysis framework. They view hedges as actions that reduce vulnerability and enhance the flexibility to respond if the unexpected occurs.⁵ Thus, for example, manufacturers might invest in flexible production facilities that can be profitable across a range of production volumes to compensate for unpredictable demand; military commanders might insist on overwhelming numerical superiority in order to compensate for the fog of war. Similarly, there may be many actions that reduce vulnerability to climate change or enhance society's ability to respond. However, such hedges will often be suboptimal for any given set of expectations about the future and hard to find in an analysis, and within a mindset, designed to produce optimum policies.⁶

Thus, an optimum policy is a poor foundation on which to build climatechange policy. Any such policy recommendation is vulnerable to attack from others who hold, or choose to hold, alternative expectations of the future. Even if forceful leadership combined with dramatic external events impel acceptance of one climate-change policy, any optimum policy could still be inadequate if the unexpected occurs, as it usually does.

2. A Method for Finding Robust Strategies

The key step in solving a complex problem is often asking the right question. Prediction-based policy analysis requires that we ask 'what is likely to happen in the future?' We believe that the proper question is 'what actions should we take, given that we cannot predict the future?' The answer we propose is that society should seek strategies that are robust against a wide range of plausible climate-change futures. By definition, robust strategies are insensitive to uncertainty about the future. For risk-averse policy-makers, such strategies would perform reasonably well, at least compared to the alternatives, even if confronted with surprises or catastrophes. Robust strategies may also provide a more solid basis for consensus on political action among stakeholders with different views of the future, because it would provide reasonable outcomes no matter whose view proved correct.

Clearly, robust strategies are desirable. The question is, do such strategies exist and, if so, do we have the means to find and assess them?

Over the last several years, we have developed a set of computational methods for decision-making under conditions of deep uncertainty that are well-suited to find and assess robust strategies for climate change (Lempert et al., 1996, henceforth LSB; Lempert et al., 2000, henceforth, LSBA; Robalino and Lempert, forthcoming, henceforth RL). These methods, called exploratory modeling (Bankes, 1993; Bankes and Gillogly, 1994), are designed to exploit the qualitatively new capabilities of modern computers, in particular, large quantities of inexpensive memory; fast, networked processors; and powerful visualization tools. The basic idea is to 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 credible distributions are available, one can lay them across the scenarios and thus calculate expected values for various strategies, value of information, and the like. In the past, the practical limitations imposed by relatively scarce computer cycles and memory forced analysts to rely on a single system model, a single set of subjective probabilities, and a single loss function as the sole representation of the future, whether or not such best-estimates were a good summary of our knowledge. The qualitatively new capabilities of modern computers now enables us to run the traditional analytic machinery of prediction-based policy analysis numerous times. We can utilize this information by compiling the expected *regret* of alternative strategies, defined as the difference between the expected performance of the strategy calculated with a particular set of priors and the performance of the optimal policy under perfect information (for a particular loss function). We then use computer search routines to look for: (i) strategies that are robust across the alternative models, priors, and loss functions (that is, their expected regret is never large compared to the alternatives); and (ii) key tradeoffs among strategies and assumptions that drive the choice among strategies. Often these searches are configured to alternatively help suggest and then to test hypothesis about robust strategies. For instance, recognizing that strategy A outperforms strategy B in some plausible futures but does poorly in others, an analyst might guess that a particular mixture of A and B might perform reasonably well across all futures, and could test this hypothesis by running the new strategy against the set of plausible futures.

The exploratory-modeling approach thus uses new technology to better implement several classic ideas for addressing deep uncertainty. As applied to the climate-change problem, we use exploratory modeling to work in a cost-benefit, Bayesian-decision-analysis framework, similarly to many of the classic treatments of climate-change policy (Nordhaus, 1994; Manne and Richels, 1992). However, we assume that there is deep uncertainty about the costs and benefits because the likelihood of some key outcomes is unknown and/or there are disagreements about how to value outcomes. Acknowledging this deep uncertainty allows us to address many of the critiques of the cost-benefit approach posed by advocates of alternative, often constraint-based methods such as the Tolerable Windows Approach (Toth, 1998). The concept of robust strategies is closely related to Simon's satisficing strategies (that is, strategies that do well-enough) and to Savage's (1950) idea of minimizing the maximum regret. In recent years, researchers have proposed formal treatments of some of these ideas, using Bayesian analysis (Berger, 1985) and robust control theory (Zhou et al., 1996), but these formalisms are difficult to apply in practice to the types of models often used in the analysis of climate-change policy. Stochastic optimization methods can also be useful (e.g., Gritsevskyi and Nakicenovic, 1999; Kelly and Kolstad, 1996), but are also challenging to employ. Our approach provides a systematic method, supported by numerical engines based on search algorithms, that simplifies the process of finding robust strategies for arbitrary models and priors. Generally, we report results by tracing tradeoff curves and by identifying the regions of parameter or probability space implicit in alternative decisions, consistent with the policy-region analysis of Watson and Buede (1987). Our contribution is to simplify the problem of inverting a Bayesian analysis by running search engines back across the inputs of computer simulations designed to do the forward problem.⁷

One advantage of the exploratory modeling approach is that it allows great flexibility in the mathematical models we can use to consider potential robust strategies. Most traditional optimization techniques limit the types of feedbacks that can be included in the underlying model. Since promising candidates for robust strategies often employ complex information feedbacks, this reduces the strategies that such optimization methods can consider. Alternatively, other simulation approaches often treat such complex feedbacks by employing a 'flight simulator' approach (Holland, 1995) in which the analyst personally examines a small number

of potential paths into the future and reports on those that seem most interesting. Such analyses offer scant basis on which to extrapolate insights to the cases not considered. In contrast, exploratory modeling makes systematic arguments from simulation models, but imposes few analytic constraints on the types of models and data used in the analysis. By treating the output of simulation models as a large, multi-dimensional ensemble of scenarios, we decouple the running of the models from the analytic techniques used to explore the models' implications.

Another important benefit of our approach is that we can easily use the language of scenario-based planning within a computational framework. The scenario-planning framework is a useful one for climate-change policy because it encourages people to grapple with the unexpected, ranging from potential catastrophes to unexpected good fortune (see, for instance, McKibbon (1998) and Schwartz (1997) for a popular press example of each). Scenario planning, along with the 'constructivist' literature on risk, also stresses the importance of multiple views in transmitting and receiving information about risk to decision-makers and stakeholders (van Asselt and Rotmans, 1997) and provides a language for communicating to decision-makers both the idea and design of robust strategies (van der Heijden, 1996). In a world of extreme uncertainty, scenario planning speaks of robust strategies comprised of *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).

Accordingly, our approach is one member of the emerging school of what we might call computational, scenario-based approaches (Morgan and Dowlatabadi, 1996; Rotmans and de Vries, 1997), in which analysts use simulation models to construct different scenarios and, rather than aggregate the results using a probabilistic weighting, argue from comparisons of fundamentally different, alternative cases. For instance, Morgan and Dowlatabadi (1996) show that climate-change strategies can be robust against the different values and objective functions held by different stakeholders. They use a small number of runs of their ICAM-2 model to demonstrate that different regional decision-makers can agree on a climate-policy choice without agreeing on the decision rules which produce that choice. Yohe and Wallace (1996) have similarly found near-term strategies in a sequential-decision framework that are robust against different expectations about the climate-change future.

Van Asselt and Rotmans (1997) pursue a view of scenario-based uncertainty similar to ours. They define three different perspectives – egalitarian, hierarchist, and individualist – that reflect different expectations about the future of climate change and three different sets of preferred strategies for addressing the climate-change challenge. Using their TARGETS model, these researchers compare the three different strategies against the three different world views and examine in depth the three 'utopias' that result when the expectations match the strategies and the six 'dystopias' that result when they don't.

Kalagnanam et al. (1998) in contrast, exploit the methodological idea of placing model outputs in a database to facilitate subsequent exploration. Specifically, they create a large database of model outputs chosen by Monte Carlo sampling and then employ a reweighting scheme to recalculate expected values of various outcomes with new probability distributions without having to recalculate the model.⁸ Casman et al. (1999) have also recently suggested using multiple models and elicited priors as 'test beds' for examining the relative robustness of alternative strategies.

Each of these approaches, in varying degrees, recognizes that there is no optimum policy that can adequately capture society's collective knowledge and views about climate change. Each addresses this fact by comparing policies against multiple views of the future.

3. Are Robust Strategies Possible?

The analytic ability to find robust strategies does not, of course, guarantee that they exist. Our exploratory modeling work on robust strategies argues, however, that adaptive decision-strategies can be a robust response to climate change. Such strategies are designed with the expectation that they will be adjusted in the future based on observations of changes in the climate and economic systems. In LSB we compared the performance of a very simple adaptive-decision strategy with that of two, commonly proposed, static alternatives, 'Do-a-Little' and 'Emissions-Stabilization'. As the names imply, the 'Do-a-Little' policy has no near-term emissions reductions and is similar to that advocated by many opponents of the commitments negotiated at the Conference of Parties in Kyoto in December 1997. The 'Emissions-Stabilization' policy returns and holds global emissions close to their 1990 levels through the mid-21st century. It is similar to the policies proposed by many advocates of the Kyoto agreement. The simple adaptive-decision strategy begins with a near-term emissions reductions intermediate between 'Do-a-Little' and 'Emissions-Stabilization', and can subsequently switch to a rate equivalent to or faster than 'Emissions-Stabilization' if damages rise or abatement costs drop sufficiently quickly.9

We compared the performance of these three strategies, using the present value of net costs and benefits as the measure, against a broad range of plausible futures, including those in which damages due to climate change turn out to be very large, those in which damages are very small, those in which technological innovation radically reduces the cost of greenhouse-gas-emissions abatement, and those in which it does not. We compared the expected value of these strategies for a wide range of expectations about the likelihood of these alternative futures, and found that even a very simple adaptive-decision strategy on average significantly outperforms either of the best-estimate policies unless our predictions of the future are highly accurate – on the order of 95%.

394

This result is not particularly surprising. The 'Do-a-Little' and 'Emissions-Stabilization' policies perform well if their underlying assumptions turn out to be valid, but they fail severely in those cases where their assumptions turn out to be wrong. The adaptive-decision strategy can make midcourse corrections and avoid significant errors. The proponents of prediction-based policy analysis might object that, in practice, most policies are adaptive and that analysts can mimic this by re-estimating their optimum policies every few years as new data comes in and as people reevaluate their subjective probabilities. But this begs the question. Our analytic tools ought to focus to the extent possible on the decision-makers' real challenge – how to begin and manage a long-term process for responding to a difficult and highly uncertain threat.

Nominally, the focus of the international political process addressing climate change is on binding targets and timetables for the reduction of near-term greenhouse-gas emissions. Prediction-based policy analysis, with its intrinsic assumption that there is some optimum level of near-term reductions, is naturally supportive of this emphasis. However, the actual political activity in this area more closely resembles an evolving set of actions designed to shape the future political landscape and influence private-sector investments, rather than any firm consensus about the optimum level of emissions reductions. International negotiators often appear to regard the continuation of the process as more important than the particular results at each stage. An incremental approach is probably quite sensible given the political constraints, but it raises the broader question of whether or not the current potpourri of actions is sufficiently robust. That is, are the world's nations taking a combination of actions that will avoid major failures no matter what future comes our way? Given that policy actions will change over time, is society currently placing too little effort in some areas and too much in others? And given that there is at present no way to determine the correct level of emissions reductions, what ought to be the goals of climate-change policy?

An analytic approach that searches for robust, adaptive-decision strategies offers a framework to address such questions. As a first step, our work suggests that climate change ought to be viewed more as a contingency problem than an optimization problem. Either society will have to make very large reductions in greenhouse-gas emissions over the course of the next century, or it will not. Since society does not yet know which future will happen, it needs to prepare for both. Thus, the most important near-term goals of climate-change policy may be to: (i) reach consensus on the observations that will indicate a need for drastic action to curtail greenhouse-gas emissions, (ii) take actions that will make large future reductions more feasible and less costly if, in fact, they are needed, and (iii) do this in such a way as to avoid over-allocating resources to a problem that turns out to be minor or preparing insufficiently for what we discover is an emerging catastrophe.

In our recent work we have made initial forays into addressing these issues. In LSBA we examine the performance of adaptive-decision strategies in the face of climate variability. This is a key question because variability may mask adverse

trends until it is too late to act or, conversely, mislead society into taking actions that are too strong. We find that in most cases adaptive-decision strategies are still robust even with quite large levels of climate variability. We have also examined the proper mix of carrots, in the form of technology incentives, as well as sticks, in the form of carbon taxes, that ought to make up an adaptive decision-strategy (RL). Much of the theoretical climate-change-policy literature suggests that carbon taxes and/or emissions trading are the most efficient climate-change policies. We find that if decision-makers hold only moderate expectations that future cost reductions due to increasing returns to scale and that heterogeneity in technology preferences among economic actors will be important, technology incentives should be a component of a robust climate-change strategy.

An important argument emerges from this very preliminary work – that policymakers may have more flexibility in designing a robust, adaptive-decision strategy for climate change than they have in choosing policies aimed at meeting some fixed emissions-reduction target. In our work on climate variability (LSBA), we consider a large number of alternative adaptive-decision strategies that differ in the rate of near-term emissions reductions, the confidence they demand that observed damage trends are real before responding to them, and the aggressiveness with which they respond to potential innovations that could reduce the cost of greenhouse-gas reductions. We find a set of robust, adaptive-decision strategies as a function of expectations about the future of climate change. For each set of expectations of the future, we find the adaptive-decision strategy that is close to the optimum for that set of expectations and also performs well over a wide variety of other expectations.

Two interesting patterns emerge as one examines this set of robust 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.

This is not to say that the targets and timetables approach embodied in Kyoto is necessarily wrong. It is of course a political judgment as to whether or not the most controversial elements ought to be at the center or the periphery of diplomatic negotiations and whether targets and timetables are the most appropriate banner under which to marshal support for the near-term actions that will lead to any necessary long-term emissions reductions. The role of policy analysis, however, is

to ensure that the decision space is clear. The search for robust, adaptive-decision strategies suggests that a robust response to climate change should employ a number of different types of government and private sector actions, from technology development to emissions trading, and that there is a wide variety of reasonable (and a wider variety of unreasonable) combinations of such actions. It is thus interesting to note that perhaps the most important policy conclusion to emerge to date from the prediction-based approach (Nordhaus, 1998) is not any agreement about the appropriate level of emissions reductions – the nominal subject of the work – but the argument that the most robust way to achieve any given emissions target is to allow for flexibility in the timing and location of reductions through some type of market-based approach. Our work supports this view of the importance of flexible mechanisms. More broadly, it suggests that a climate-change strategy that succeeded in: (1) establishing the physical and institutional capability to monitor the relevant climate and economic systems, (2) establishing the capability to effectively regulate greenhouse gases, and (3) encouraging the use of new emissions-reducing technologies (Jacoby et al., 1998), would be at least as successful in the long-term as a policy that succeeded in meeting particular near-term targets for reductions in greenhouse-gas emissions.

4. Implications

The practitioners of scenario-based planning (Schwartz, 1996) generally oppose efforts to assign probabilities to different scenarios of the future. Scenario-planners regard scenarios as a handful of plausible stories about the future designed to illustrate the alternative consequences of key uncertainties. These scenarios are intended to help groups of decision-makers, such as the senior management of a corporation, internalize the fact that the future may be quite different from a simple extrapolation of the present. Motivated by these scenarios, decision-makers can better craft policies that hedge against the risks they face. Assigning probabilities to scenarios would interfere with this process because people are likely to assign a high probability to those scenarios with which they are most comfortable and a low probability to those scenarios they find most troubling. If they are able to discount the likelihood of uncomfortable futures, people tend to ignore them when making their plans.

Scenario planning thus provides a useful language for focusing the attention and imagination of decision-makers, considered as individuals and as groups, on a highly uncertain future. But scenario planning generally lacks the great advantage of traditional prediction-based policy analysis – an ability to bring to bear a great deal of quantitative information to systematically adjudicate among different policy alternatives. We believe that the qualitatively new capabilities provided by today's computer technology have now made possible a synthesis between the more natural language of scenario-based planning and the ability of probability-

based decision analysis to systematically use quantitative information to compare policy options. This view of the problem – looking for robust strategies rather than optimum policies – has implications for the users of simulation models as well as their suppliers, because it affects the demands for policy-relevant information from the scientific community and the types of questions policy analysts can help policymakers ask. For instance, here are two examples of questions that seem amenable to systematic, quantitative analysis, but which are generally ignored because they fit poorly with the framework of prediction-based policy analysis.

First, what are reasonable and unreasonable combinations of shaping and hedging actions for a climate-change policy? Most climate-change policy and policy analysis concentrates on near-term emissions reductions and the steps, such as emissions-trading, needed to achieve these reductions. But are these reductions intended as a hedging or a shaping action? Most analysis and policy debate suggests the former, that the purpose of near-term reductions is to reduce atmospheric concentrations of greenhouse gases and thus hedge against the possibility of serious impacts from climate change. But if climate change proves to be extreme, these reductions will have little benefit. However, such reductions may be useful as shaping actions, intended to make very large reductions easier in the future if those prove necessary. The primary benefit of Kyoto may be to spur investment in emissions-reducing technologies and initial experimentation with an international emissions-trading scheme that may take decades to perfect. If such is the purpose, or at least a significant potential benefit of current policies, then the question of optimal emissions paths is much less helpful than, for instance: (i) information on the range of emissions-reductions paths with which emissions trading institutions may someday have to contend, and (ii) analysis of the economic, political, and other shocks that might derail such institutions before they have a chance to become established.

Second, how do we know whether or not our climate policies are working? What are the signposts that would suggest we have to change our strategy? There is much literature on fingerprints of climate change and whether or not particular already-observed trends are due to anthropogenic sources. But few have bothered to estimate which observations over the next decade or so should cause us to validate or invalidate our current climate-change strategy. In particular, literatures as diverse as scenario-based planning and control theory both make clear there is a close connection between available signposts and the choice of strategy. In other words, in a world of deep uncertainty, one chooses a strategy in large part based on what reliable observations one can make to re-adjust that strategy over time. To date, climate policy rests on observations of emissions of greenhouse gases, which are probably a very poor signpost if for no other reason that, in the short-term at least, variations in emissions are affected far less by climate policy than they are by factors such as variations in economic growth (see, for instance, Schelling, 1992). The climate-change-policy community could do policy-makers a

great service by examining signposts that might provide a better basis for building near-term climate-change policies.

Questions such as these cry out for a quantitative assessment, and require key inputs from the scientific community. But such questions are hard to address with a prediction-based policy analysis predicated on answering 'what is likely to happen in the future?' Society should pursue robust strategies that will work reasonably well no matter what the future brings. Fortunately, new computer technology has provided analysts the means to systematically propose and evaluate such approaches to the threat of climate change.

Notes

¹ Or many of us, if the juggernaut of medical science sufficiently slows the process of aging within our lifetimes.

 2 As adjusted by any risk-aversion on the part of the decision-maker.

 3 Our criticisms are similar to those found in IPCC (1996), though we believe we also offer some innovative remedies.

⁴ To be sure, empirical evidence and scientific arguments can narrow the different expectations among different groups. For instance, the scientific community might generate a consensus about the subjective probability distribution for some physical parameter, such as climate sensitivity, sufficient to convince any reasonable stakeholder. But there are many key parameters for which this is just not possible, for instance, the likelihood that new technology will substantially lower the cost of greenhouse-gas abatement.

⁵ This language is common in the qualitative literature on scenario-based planning (Dewar, 1993) and business strategy. Finance theory provides a quantitative description of optimum hedges that reduce vulnerability given certain assumptions (e.g., liquidity, known probabilities of price movements, etc.). These hedges are not necessarily optimum, nor prudent, in cases where those assumptions do not hold.

 6 This is the reason for the seemingly paradoxical unimportance of uncertainty in many analytic discussions of climate-change policy. Often we find that the optimum policy is little affected when uncertainty is subsequently added to the analysis and thus, uncertainty may seem unimportant. But such analyses do not provide any language to discuss the actual response of decision-makers to conditions of deep uncertainty – adopt strategies designed to work reasonably well against a wide range of plausible futures.

⁷ It is useful to compare exploratory modeling with sensitivity analysis. In the latter we first assume our models and data are accurate, calculate an optimum policy, and then ask how sensitive that policy is to changes in our assumptions. Exploratory modeling asks the question in reverse order. It first assumes that our models cannot predict the future nor the likelihood of different events, then provides the means to look for strategies that perform well across a broad range of plausible futures.

⁸ This use of Monte Carlo sampling is consistent with the concept of exploratory modeling, and we have used it in our work (Robalino and Lempert, forthcoming). However, it is important that one save the results of individual cases to use in further analysis rather than lose this information by saving only the mean and variance of the outputs over the sample. In addition, Monte Carlo sampling is only one way to search over an ensemble of scenarios. Often in an exploratory modeling analysis we want to use a search to find particular cases. For instance, we might use a genetic algorithm search (e.g., Miller, 1998) to find those scenarios in which a particular strategy performs best and those in which it performs worst.

⁹ More complex adaptive-decision strategies are certainly possible. We consider such strategies in the work discussed below.

References

- Bankes, S. C.: 1993, 'Exploratory Modeling for Policy Analysis', *Oper. Res.* **41**, 435–449 (also published as RAND RP-211).
- Bankes, S. C. and Gillogly, J.: 1994, 'Exploratory Modeling: Search Through Spaces of Computational Experiments', in Sebald, A. V. and Fogen, L. J. (eds.), *The Proceedings of the Third Annual Conference on Evolutionary Programming*, World Scientific, River Edge, NJ (also published as RAND RP-345).
- Berger, J. O.: 1985, Statistical Decision Theory and Bayesian Analysis, Springer-Verlag, Berlin, Heidelberg. See Section 4.7 on Bayesian Robustness.
- Casman, E. A., Morgan, M. G., and Dowlatabadi, H.: 1999, 'Mixed Levels of Uncertainty in Complex Policy Models', *Risk Anal.* 19, 33–42.
- Dewar, J. A., Builder, C. H., Hix, W. M., and Levin, M. H.: 1993, Assumption Based Planning: A Planning Tool for Very Uncertain Times, RAND, Santa Monica, CA, MR-114-A.
- Gritsevskyi, A. and Nakicenovic, N.: 1999, 'Modeling Uncertainty of Induced Technological Change', Paper presented at the Pew Center Workshop on Economics and Integrated Assessment of Climate Change, 21–22 July, Washington, D.C.
- Holland, J. H.: 1995, *Hidden Order: How Adaptation Builds Complexity*, Addison-Wesley, Menlo Park, CA.
- IPCC (Intergovernmental Panel on Climate Change): 1996, in Bruce, J. P., Lee, H., and Haites, E. F. (eds.), *Climate Change 1995: Economic and Social Dimensions of Climate Change*, Cambridge University Press, Cambridge.
- Jacoby, H. D., Prinn, R. G., and Schmalensee, R.: 1998, 'Kyoto's Unfinished Business', Foreign Affairs 77, 54–66.
- Johansson, T. B., Kelly, H., Reddy, A. K. N., and Williams, R. H.: 1993, *Renewable Energy: Sources for Fuels and Electricity*, Island Press, Washington, D.C.; Covelo, CA.
- Kalagnanam, J. R., Kandlikar, M., and Linville, C.: 1998, 'Importance Ranking of Model Uncertainties: A Robust Approach', *Risk Anal.*
- Kelly, D. L. and Kolstad, C. D.: 1996, 'Tracking the Climate Change Footprint: Stochastic Learning about Climate Change', Economics Working Paper 3-96R, University of California Santa Barbara.
- Lempert, R. J., Schlesinger, M. E., and Bankes, S. C.: 1996, 'When We Don't Know the Costs or the Benefits: Adaptive Strategies for Abating Climate Change', *Clim. Change* 33, 235–274.
- Lempert, R. J., Schlesinger, M. E., Bankes, S. C., and Andronova, N. G.: 2000, 'The Impacts of Variability on Near-Term Policy Choices and the Value of Information', *Clim. Change* 45, 129– 161.
- Manne, A. S. and Richels, R. G.: 1992, Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emissions Limits, MIT Press, Cambridge, Mass.
- McKibben, B.: May 1998, 'A Special Moment in History', Atlantic Monthly 281, 55-78.
- Miller, J. H.: 1998, 'Active Nonlinear Tests (ANTs) of Complex Simulation Models', *Manage. Sci.* 44, 820–830.
- Morgan, M. G. and Dowlatabadi, H.: 1996, 'Learning from Integrated Assessments of Climate Change', *Clim. Change* 34, 337–368.
- Morgan, M. G. and Keith, D.: 1995, 'Subjective Judgments by Climate Experts', *Environ. Sci. Technol.* 29, 468A–476A.

- Morgan M. G. and Henrion, M.: 1990, Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis, Cambridge University Press, Cambridge, p. 332.
- National Research Council: 1997, 'Modeling and Simulation', Vol. 9 of Technology for the United States Navy and Marine Corp: 2000–2035, National Academy Press, Washington, D.C.
- Nordhaus, W.: 1998, 'Integrated Assessment Modeling of Climate Change', in Nordhaus, W. (ed.), *Economics and Policy Issues in Climate Change*, Resources for the Future, Washington, D.C.
- Nordhaus, W. D.: 1994, *Managing the Global Commons: The Economics of Global Change*, MIT Press, Cambridge, Mass., p. 213.
- Repetto, R. and Austin, D.: 1997, *The Costs of Climate Protection: A Guide for the Perplexed*, World Resources Institute, Washington, D.C.
- Robalino, D. A. and Lempert, R. J.: 2000, 'Carrots and Sticks for New Technology: Crafting Greenhouse Gas Reductions Policies for a Heterogeneous and Uncertain World', *Integrated Assess*. (forthcoming).
- Rotmans, J. and de Vries, H. J. M.: 1997, *Perspectives on Global Change: The TARGETS Approach*, Cambridge University Press, Cambridge.

Savage, L. J.: 1950, The Foundations of Statistics, John Wiley and Sons, New York.

- Schelling, T. C.: 1992, 'Some Economics of Global Warming', Amer. Econ. Rev. 82 (March), 1-14.
- Schwartz, P.: 1996, The Art of the Long View, Currency Doubleday, New York, p. 272.

Schwartz, P.: 1997, 'The Long Boom', Wired 5.07, 115-129, 168-173.

- Sherden, W. A.: 1998, *The Fortune Sellers: The Big Business of Buying and Selling Predictions*, John Wiley and Sons, New York.
- Toth, F. L.: 1998, Cost-Benefit Analysis of Climate Change: The Broader Perspective, Birkhauser Verlag, Berlin, p. 160.
- van Asselt, M. B. A. and Rotmans, J.: 1997, 'Uncertainties in Perspective', in Rotmans, J. and de Vries, B. (eds.), *Perspectives on Global Change: The TARGETS Approach*, Cambridge University Press, Cambridge.
- van der Heijden, K.: 1996, *Scenarios: The Art of Strategic Conversation*, John Wiley and Sons, New York, p. 305.
- Watson, S. and Buede, D.: 1987, Decision Synthesis, Cambridge University Press, Cambridge.
- Yohe, G. and Wallace, R.: 1996, 'Near Term Mitigation Policy for Global Change under Uncertainty: Minimizing the Expected Cost of Meeting Unknown Concentration Thresholds', *Environ. Model.* Assess., 47–57.
- Zhou, K., Doyle, J., and Glover, K.: 1996, *Robust and Optimum Control Theory*, Prentice-Hall, New Jersey.

ROBERT J. LEMPERT

RAND, 1700 Main St., Santa Monica, CA 90407-2138, U.S.A.

Department of Atmospheric Sciences, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A. MICHAEL E. SCHLESINGER