Characterizing and Dealing With Uncertainty: Insights from the Integrated Assessment of Climate Change[1]

M. GRANGER MORGAN

Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA

ABSTRACT

After posing four-broad questions about uncertainty and climate change, Part 1 of this paper provides a review of basic ideas about uncertainty and its treatment in quantitative policy analysis. Part 2 reports very briefly on a series of expert elicitations of climate experts which the author and his colleagues have conducted. The final portion of the paper uses integrated assessment of climate change as a vehicle to explore some limitations to conventional policy analysis and the treatment of uncertainty. Because the climate problem is global in scope, involving many societies, and because it will involve large changes that unfold on a time-scale of several centuries, many standard analytic methods for policy analysis cannot be appropriately applied. Several such difficulties are identified and explored. Strategies for addressing a few of them are discussed. The paper closes by offering answers to the four opening questions.

Keywords: uncertainty, climate change, expert elicitation, integrated assessment.

1. SOME QUESTIONS

We begin with a deceptively simple question:

> What is the role of uncertainty in the assessment of future climate change and its likely ecological, social and economic implications?

Many participants in the discussion of climate change are of two minds when they try to answer.

We can frame the climate change integrated assessment problem as that of providing accurate solutions to a set of coupled models such as those illustrated in Figure 1. At least some who adopt this framing also argue that significant policy interventions are not warranted until we can obtain accurate predictive solutions to such models. We might term this . . . the "rational social decision making" formulation of the climate problem. To those of us who come from technical backgrounds, this formulation – "first figure out how the system works, and only then take action" – has some clear appeal. However, I believe that we would be hard pressed to find examples of any major human decisions which have been actually made this way. Decision to go to war, sue for peace, found a city, create a company, get married, start a family, or emigrate to another land, have always involved a leap of faith based on fragmentary information.

In decision-analytic terms, the formulation "no action before clear and complete knowledge (i.e., avoid false positives at all costs)," may be fine as the first-order criterion to choose which papers to publish in peer-reviewed scientific journals. However, it carries the implicit assumption that the costs of taking no action are small. In many fields of public decision making – such as national security, public health, and environmental protection – that may not be a safe assumption.

Which brings us to a second perspective. If we conclude that there is enough evidence now about human impacts on climate to suggest a reasonable basis for concern, and we believe that there is no way that we will be able to adequately understand and model the set of coupled processes displayed in Figure 1 (at least on the time-scale of the geophysical experiment we are currently running), then the climate problem is fundamentally a problem of: (1) producing a shared conviction among the peoples and leaders of the large industrialized and industrializing nations that anticipatory actions are needed; and (2) devising robust adaptive strategies for dealing with an inherently uncertain future. I adopt this latter view.
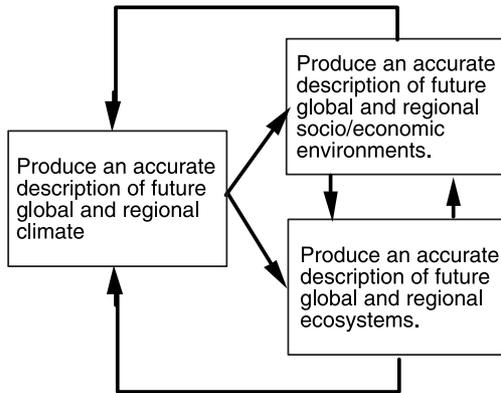

Fig. 1. The problem of integrated assessment of the climate problem is sometimes framed as the problem of providing accurate solutions to a set of coupled models such as these.

Thus, in considering uncertainty in the case of climate assessment, I believe we should focus primarily on four questions:

1. How can we use uncertainty analysis, bounding analysis, and similar methods, to limit the range of possible "answers" even if perfect understanding and predictive capability cannot be achieved over the next few decades?
2. How can we develop a clearer indication of what we can expect to know, and what we will probably not be able to know, about climate change and its likely impacts on the time-scale of the next few decades?
3. How can we use scientific knowledge and policy analysis to develop strategies for social and individual behavior and choice that are likely to be robust in an inherently uncertain world?
4. As we elaborate our characterization of uncertainty, and what we cannot expect to know, how can we do it in such a way as not to provide aid to those who argue that no action is warranted until all uncertainties have been resolved?

I will return to these four questions to offer my personal "answers" at the end of this paper.

2. BASIC IDEAS ABOUT UNCERTAINTY

Uncertainty is an integral part of most large problems in science, technology and public policy, of which climate change is a prime example. The quantitative characterization and treatment of uncertainty has been a central preoccupation of the fields of decision and risk analysis since their inception. Raiffa and Schlaifer [8], Howard and Matheson [9], Rasmussen and colleagues (USNUREG, 1975) [10], Keeney [11], von Winterfeldt and Edwards [12], Morgan and Henrion [13], Evans et al. [14], Paté-Cornell [15], Callahan [16] and many others have contributed to the systematic exploration of the problems involved. They have created a

solid theoretical framework for addressing uncertainty and have developed and demonstrated a range of practical tools for its characterization, and its incorporation into quantitative risk and other types of policy analysis.

Most of this literature is focused on how to characterize and analyze uncertainty about the value of specific coefficients that arise within a model or the value of input variables to that model. Here are two climate-related examples:

- Climate sensitivity – If we double the concentration of carbon dioxide in the earth's atmosphere, how much will the average temperature of the earth increase?
- Future energy prices – What will be the bus-bar price of electricity from conventional coal fired power plants in Western Pennsylvania in January of 2020?

Morgan and Henrion [13] provide a classification of a variety of sources of such coefficient or parameter uncertainty, and discuss how each should appropriately be described and analyzed.

The standard way to describe such uncertainty is with a probability density function (PDF) or its integral, the cumulative distribution function (CDF). This can be done using a frequentist approach when adequate data are available. More typically, data are incomplete, and it is necessary to seek the considered judgment of experts using a subjectivist approach in which probability distributions become statements of "degree of belief." In Part 2, I discuss the problem of eliciting full subjective probabilistic judgments from experts and present several climate-specific examples,

The use of quantitative statements of uncertainty has a number of clear advantages. However, there remain people who argue that it is sufficient, or even preferable [17], to describe uncertainty in terms of "probability words" (e.g., words such as "likely" and "unlikely"). Typically the folks who advance such arguments also adopt a "frequentist" perspective on probability, rejecting Bayesian "subjectivist" interpretations of probability as a statement about degree of belief [13].

Qualitative language used without at least some quantification is inadequate because: (1) the same words can mean very different things to different people; (2) the same words can mean very different things to the same person in different contexts; (3) important differences in experts' judgments about mechanisms (functional relationships), and about how well key coefficients are known, can be easily masked in qualitative discussions.

Figure 2, which summarizes work by Wallsten et al. [18], illustrates the range of meaning that common probability words can take on for different individuals, absent any specific problem context. Note the enormous ranges associated with some words and phrases such as "good chance" and "possible," and the large variation in the interpretation of the words by different subjects, and the significant overlap between a number of words.
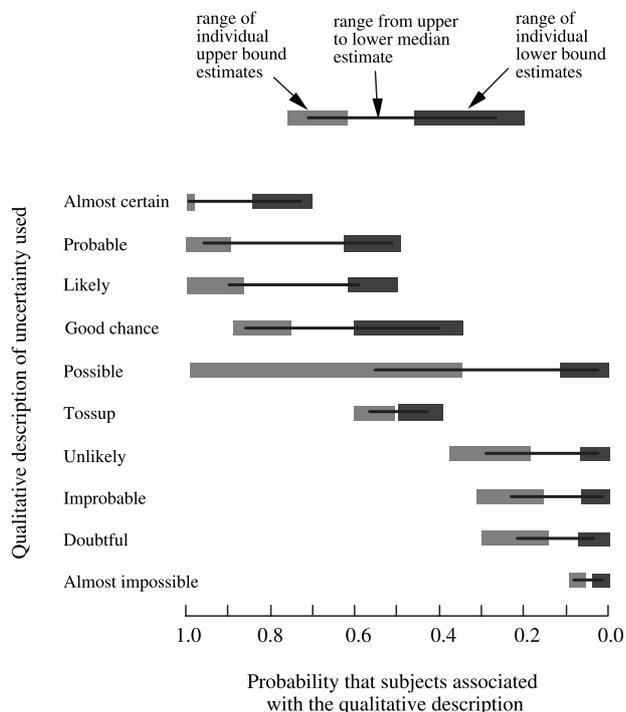
Fig. 2. Range of quantitative values that Wallsten et al. [18] found that people associate with common probability words absent any specific problem context. Note the enormous ranges associated with some words and phrases such as "good chance" and "possible," the large variation in the interpretation of the words by different subjects, and the significant overlap between a number of words.
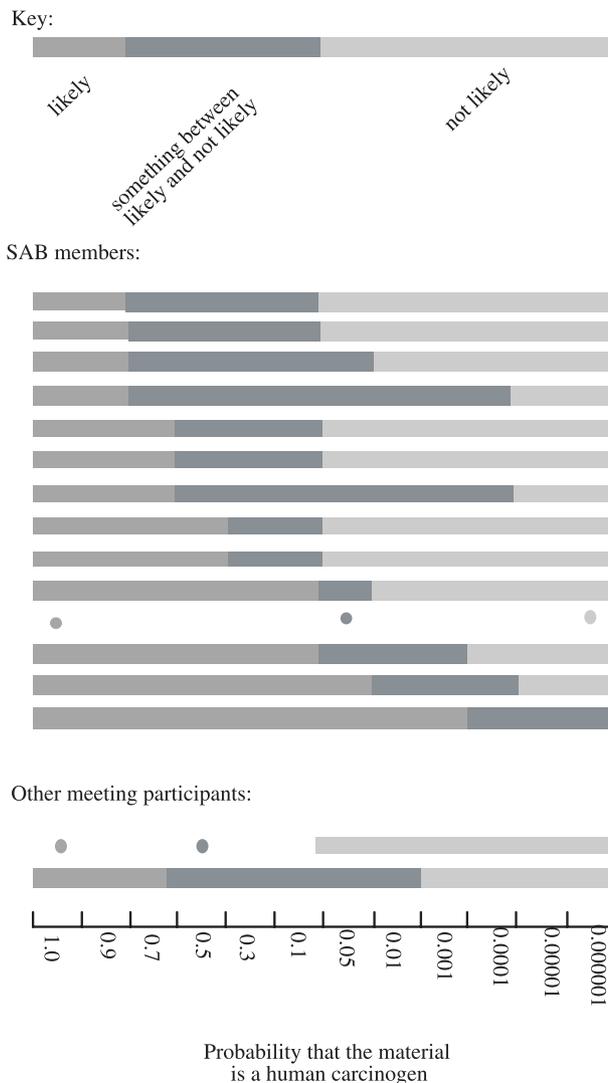

Fig. 3. Results obtained when members of the executive committee of the US EPA Science Advisory Board were asked to assign quantitative values to the probability words "likely," "not likely" and to "something in between likely and not likely" in the context of the possible carcinogenicity of chemicals [5]. Note that, even in this relatively small and expert group, the minimum probability associated with the word "likely" spans four orders of magnitude, and that the maximum probability associated with the word "not likely" spans more than five orders of magnitude. Indeed, even in this small sample, there is an actual *overlap* of the probability associated with the word "likely" and that associated with the word "unlikely." Note that two respondents provided only point values, not ranges.

To illustrate the problem more specifically, in 1997 I asked the members of the Executive Committee of the EPA Science Advisory Board to give me their quantitative interpretations of some of the words being proposed for use in EPA's new cancer guidelines [19]. Specifically I asked them to indicate what range of numerical probabilities they would assign to "likely," "not likely" and to "something in between likely and not likely" [5]. Figure 3 summarizes the responses I received. Note that, even in this relatively small and expert group, the minimum probability associated with the word "likely" spans four orders of magnitude, and that the maximum probability associated with the word "not likely" spans more than five orders of magnitude. Indeed, even in this small sample, there is an actual *overlap* of the probability associated with the word "likely" and that associated with the word "unlikely!"

Clearly, without at least some quantification, such qualitative descriptions of uncertainty convey little, if any, useful information. The climate assessment community appears to be well along in learning this lesson. Moss and Schneider [20] have worked to get a better treatment of uncertainty incorporated in the current round of IPCC. Progress is uneven, but awareness is growing. In the summary for policy makers recently released by IPCC WG 1 we read:

> ...the following words have been used...to indicate judgmental estimates of confidence: virtually certain (greater than 99% chance that the result is true); very likely (90–99% chance); likely (66–90%); medium likelihood (33–66%); unlikely (10–33% chance); very unlikely (1–10% chance); exceptionally unlikely (less than 1% chance)...

There is also a notably increased focus on discussions of uncertainty, and on including uncertainty in graphical results. The US National Assessment Synthesis Team [21] also gave quantitative definitions to five probability words and tried to use them consistently throughout their overview report.

A second, often more important type of uncertainty, is uncertainty about model form. By this I mean uncertainty about how the physical, biological and social world works. Here again we can illustrate with two climate related examples:

- Amount and distribution of precipitation – If the climate warms as a result of doubling the atmospheric concentration of CO_2, which physical processes will control the amount of precipitation at middle latitudes? Will the result be an increase or a decrease?
- Predicting Chinese climate policy – In selecting a climate policy in the latter part of the century, will China choose the one that gives the best expected net present value, the one that gives the lowest maximum possible loss, or some other option?

Model uncertainties can arise because: (1) we don't know the basic science that underlies the processes being modeled; (2) the processes involved are sufficiently complex that we are unable to model their behavior from first principles, and do not have enough information to confidently build macroscopic models; (3) the system is changing in unknown ways over time. Examples of this last include multiple states (not all states understood, or transition rules between states not understood); structural change; statistical non-stationary; and chaotic behavior.

There are a variety of strategies that can be used to deal with model uncertainty. If you can spell out all the alternative possible model forms then you can analyze each separately, and combine the results (Fig. 4a). If you can't assess the probabilities, you can at least try to bound the range of possible outcomes. Similarly, if you cannot spell out all the possible models you may still be able to bound the range of outcomes with order-of-magnitude or similar methods (Fig. 4b).

Until recently there had been little practical progress in dealing with model uncertainty, but now there are several
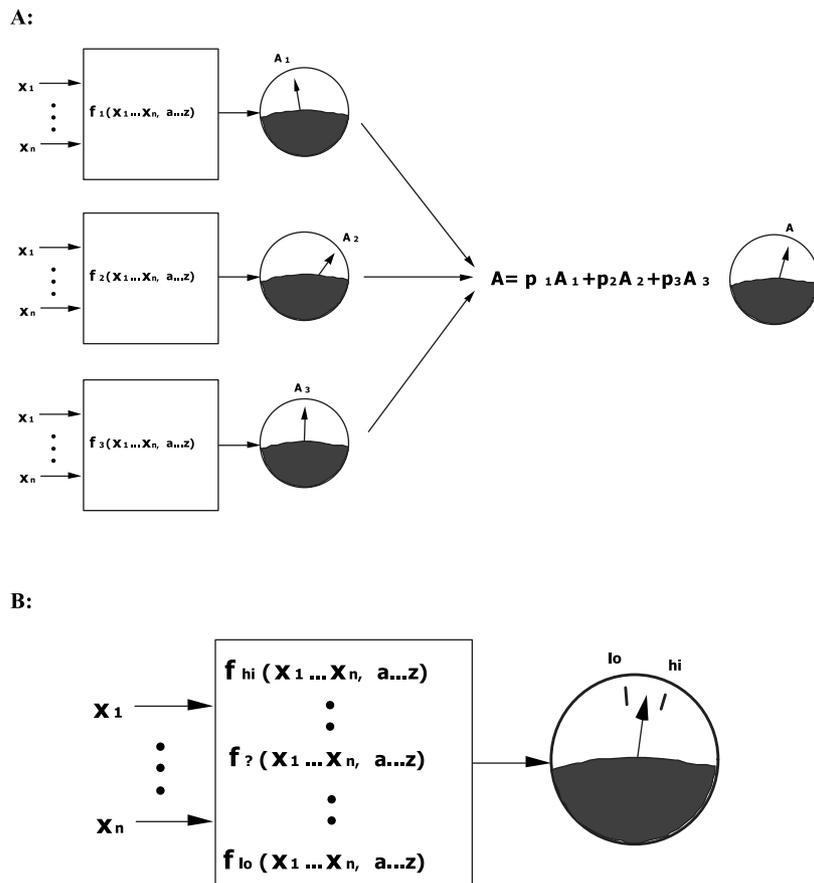

Fig. 4. There are a variety of strategies that can be used to deal with model uncertainty. If you can spell out all the alternative possible model forms, then you can analyze each separately, and combine the results (**A**, above). If you can't assess the probabilities, you can at least try to bound the range of possible outcomes. Similarly, if you cannot spell out all the possible models, you may still be able to bound the range of outcomes with order-of-magnitude or similar methods (**B**, below).

examples of serious applied efforts. John Evans and his colleagues at the Harvard School of Public Health [14] have performed a series of analyses to explore the possible mechanisms that underlie the toxicity of environmental pollutants. Alan Cornell and others have explored model uncertainty in the context of seismic risk [22]. Hadi Dowlatabadi and I, along with several colleagues at Carnegie Mellon, have worked to incorporate model uncertainty into an integrated assessment of climate change [6]. I will return to this work in Part 4.

3. CLIMATE-RELATED EXAMPLES OF EXPERT ELICITATION

Eliciting full subjective probabilistic judgments requires careful preparation and execution. Developing and testing an appropriate interview protocol typically takes several months. For any moderately detailed topic, each interview is likely to require several hours. When addressing complex, scientifically subtle questions of the sorts involved in climate change and impact assessment, there are no satisfactory short cuts. Attempts to simplify and speed up the process almost always lead to shoddy results due to factors such as: (1) imprecisely specified questions; (2) inadequate attention to overconfidence; or (3) inadequate attention to other effects of cognitive heuristics such as "availability," "anchoring and adjustment" and "representativeness" [23].

Developing a precise question can often be more complex than one might think. For example the price of gas in 2010 is not a precise definition of an uncertain variable. Are we talking regular or high octane, wholesale or retail, in the US or Netherlands? If in the US, what state? What time of year? And so on.

There is a large literature, summarized in [13], which clearly demonstrates that both ordinary people and experts are systematically overconfident about the precision with which they know the things they know. This can have the impact of producing elicited distributions which are much narrower than they should be.

When ordinary people or experts make judgments about uncertain events, such as numbers of deaths from chance events, they use simple mental rules of thumb called "cognitive heuristics." In many day-to-day circumstances, these serve us very well, but in some instances they can lead to bias – such as over confidence – in the judgments we make. This can be a problem for laypeople and for experts. Space limitations preclude a discussion here of these issues. A good general introduction can be found in [29]. Details, as they apply to expert elicitation can be found in [13].

Since 1994 my colleagues and I have conducted two large-scale projects of expert elicitations [1, 7]. Others who conducted (much simpler) elicitations include NDU [25], Stewart and Glantz [26] and Nordhaus [27].
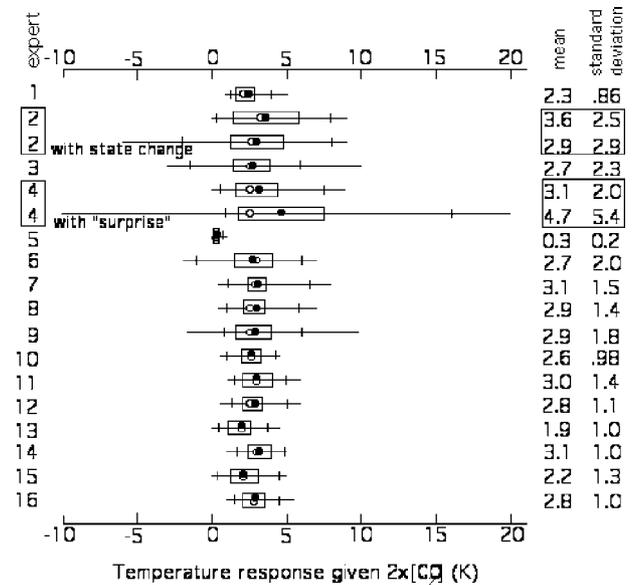

Fig. 5. Box plots of elicited probability distributions of climate sensitivity, the change in globally averaged surface temperature for a $2\times[CO_2]$ forcing. Horizontal line denotes range from minimum to maximum assessed possible values. Vertical tick marks indicate locations of lower 5 and upper 95 percentiles. Box indicates interval spanned by 50% confidence interval. Solid dot is the mean and open dot is the median. The two columns of numbers on right-hand side of the figure report values of mean and standard deviation of the distributions. From [7].

Figure 5 provides an illustration of one of the results from our elicitation of climate scientists. Additional results from our elicitations with climate scientists, as well as some results from our elicitation of forest ecosystem experts are available in [1, 7].

4. UNCERTAINTY IN THE ASSESSMENT OF THE CLIMATE PROBLEM

We turn now to a discussion of the treatment of uncertainty in the assessment of climate change. The past three decades have witnessed an explosive growth in the development and use of tools for quantitative policy analysis, among which have been tools for the characterization and treatment of uncertainty [13]. As policy analysts have turned to the consideration of climate and other problems of global change, they have found it natural to employ these standard tools without modification. However, many issues in global change involve temporal, spatial and socio-political scales that are significantly broader than those encountered in most traditional policy analyses. In such cases, the uncritical application of conventional tools can violate the assumptions on which they are based, produce silly or misleading findings [4].
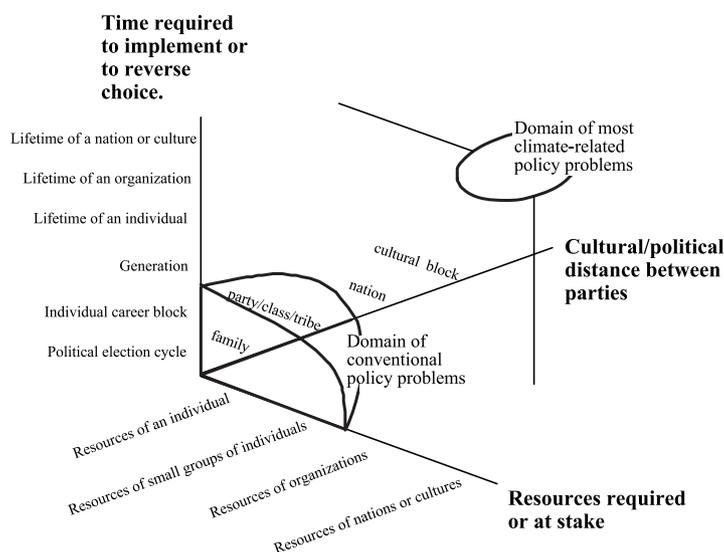
Fig. 6. Policy problems can be placed in this space and their location used to evaluate the appropriateness of applying various policy analysis tools. For any particular policy problem the axes indicate the amount of resources required or at stake; the time-scale, both to implement and to reverse the effects of the choices available; and the degree of political and cultural homogeneity of the people involved. Most tools of modern policy analysis work best when the problem lies near the origin in this space. As one moves outward from the origin, more and more of the underlying assumptions on which these tools are based break down. Many problems of global change lie far from the origin on all three axes. Global change issues may involve very large costs, are often characterized by long temporal scales and associated intergenerational equity issues, and may involve a large political and cultural distance between many different parties, and an associated lack of shared metrics. From [4].

The source of difficulty is illustrated in Figure 6. Most tools of modern quantitative policy analysis were developed to address problems that lie near the origin in this space. As one moves outward from the origin, more and more of the underlying assumptions upon which conventional tools are based begin to break down. Because many problems in global change lie far from the origin on all three dimensions, one can expect that the straightforward application of standard ideas and methods will often fail.

In other writings [4] my colleagues and I have argued that among the standard assumptions that are problematic are: (1) the assumption that there is a single public-sector decision maker who faces a single problem in the context of a single polity; (2) the assumption that the impacts involved are of manageable size and can be valued at the margin; (3) the assumption that values that are known, static, and exogenously determined, and that the decision maker should select a policy by maximizing expected utility; (4) the assumption that time preference is accurately described by conventional exponential discounting of future costs and benefits; (5) the assumption that uncertainty is modest and manageable; and (6) the assumption that for most questions of interest, the system under study can reasonably be treated as linear. Of course, not all conventional policy analysis makes all these assumptions, nor are all problems that lie near the origin in Figure 6 amenable to solution with analytical tools that make such assumptions.

Here we focus just on problems that arise in the characterization and treatment of uncertainty. We have very different levels of understanding of the various processes that determine: (1) human and natural emissions of radiatively important gases and particles; (2) climate system response to radiative forcing; (3) human and natural response to climate change. In some cases, we understand processes well enough, and/or time constants are slow enough, that we can safely model or extrapolate for centuries. But, in many cases, particularly those involving human or ecosystem responses, we do not understand processes well enough, and change can happen so quickly, that our ability to model or extrapolate is very limited.

Over the past several years, Dowlatabadi, Morgan and co-workers built a large stochastic simulation model called ICAM (for Integrated Climate Assessment Model) in the Analytica^{tm} software environment (formerly Demos). This environment provides a powerful graphic user interface which represents the model structure in the form of hierarchically organized influence diagrams. Users can explore the model by "double clicking" on various elements, moving down through the model hierarchy until they reach individual model elements, where they can observe the mathematical relationships between variables and read documentation on some of the values being used and the assumptions that have been made. Users can easily substitute alternative values or probability distributions. The model is available at http://www.hdgc.epp.cmu.edu. A demonstration copy of the Analytica^{tm} software can be obtained at http://www.lumina.com/software.

In the current version of ICAM, the world is divided into twelve regions. Time is stepped in five-year increments from 1975 to 2100. Demographic and economic processes lead to

emissions of greenhouse gases and aerosols. These modify the composition of the atmosphere, and bring about climate change. Climate change leads to various impacts which in turn can affect demographic, socio-economic and ecological processes. It is possible to make policy interventions in energy use, in emissions management, and in adapting to impacts. In some user-selected structural variants of ICAM, economic factors, climate change, and climate impacts can influence the initiation and path of these interventions.

In developing ICAM, we found that uncertainty about the appropriate functional form of different sub-models is sufficiently large, and the difficulty of constructing all plausible alternatives sufficiently great, that it is often best to report results parametrically across a set of combinations of different model structural assumptions, in much the same way that one reports the results of parametric sensitivity studies of coefficient uncertainty. For example, in an application of ICAM-2 designed to explore the probability that a specific carbon tax policy would yield net positive benefits, we found that the probability ranged from 0.15 to 0.95 for the world as a whole, depending upon the structural assumptions made. A more recent example, from a study [29] conducted of the costs of delaying mitigation activities, illustrates the effects of alternative model structures in just the energy and carbon emission control modules of ICAM-3, Table 1.

Many climate policy models are designed and solved as long-term optimization problems. In a setting with uncertainties as great as those displayed in Table 1, we doubt the utility of conventional optimization formulations. As an alternative, rather than try to search for the optimal policy, we have set out to search for robust behaviors. Just as in the model environment, real world policy makers will always face great uncertainty. They must observe the world, use what they see together with models to make forecasts, choose what they think is the best strategy at the moment, and then a few years later, repeat the entire cycle. By building simple ''decision agents'' we have been able to do something very similar within the world of the ICAM model environment. Then, across a range of alternative model worlds we run repeated stochastic simulations of the model and ask, among a range of plausible alternative behavioral strategies which our agents might adopt, which one does best in the face of the uncertainties about both coefficient values and model structures? In the case of the climate problem, a strategy that tracks and attempts to control atmospheric concentration of greenhouse gases (as opposed to emissions or temperature), using a quadratic penalty function, seems to do best. Of course, not all problems with high uncertainty will yield such a single general result. In some cases, even a recasting of the problem in terms of the behavior of autonomous agents is likely to lead to different behaviors for different combinations of model structure.

We are of course not the first to argue the importance of an adaptive approach in environmental decision making. There has been a long tradition of arguing the strength of such approaches, particularly among systems ecologists. In what is now a classic paper, Clark [29] articulated a compelling argument for the superiority of such approaches when knowledge is incomplete. Lee [30] has elaborated

Table 1. Illustration of the wide range of results that can be obtained with ICAM depending upon different structural assumptions, in this case, about the structure of the energy module and assumptions about carbon emission control. In this illustration, produced with a 1997 version of ICAM, all nations assume an equal burden of abatement by having a global carbon tax. Discounting is by a method proposed by Schelling. Other versions of ICAM yield qualitatively similar results. Table based on [28].

Model components	Model structural variants								
	M1	M2	M3	M4	M5	M6	M7	M8	M9
Are new fossil oil & gas deposits discovered?	no	yes	no	no	yes	yes	no	yes	yes
Is technical progress that uses energy affected by fuel prices and carbon taxes?	no	no	yes	no	yes	yes	yes	yes	yes
Do the costs of abatement and non-fossil energy technologies fall as users gain experience?	no	no	no	yes	no	no	yes	yes	yes
Is there a policy to transfer carbon saving technologies to non-Annex 1 countries?	no	no	no	no	no	yes	yes	no	yes
TPE BAU in 2100 (EJ) Mean	1975	2475	2250	2000	3425	2700	1450	3550	2850
TPE control in 2100 (EJ) Mean	650	650	500	750	500	500	675	750	725
CO_2 BAU 2100 (10^9TC) Mean	40	50	50	40	75	55	25	73	55
Std. Deviation	28	18	36	29	29	23	22	27	21
Mitig. Cost (%Welfare) Mean	0.23	0.44	0.14	0.12	0.48	0.33	0.05	0.23	0.17
Std. Deviation	0.45	0.23	0.23	0.22	0.28	0.12	0.07	0.12	0.11
Impact of delay (%Welfare) Mean	−0.1	0.2	−0.6	0.0	−1	−0.5	−0.1	−0.6	−0.4
Std. Deviation	1	0.3	1	0.7	1.2	0.9	0.5	0.8	0.6

Note. TPE = Total Primary Energy.

BAU = Business as Usual (no control and no intervention).

Sample size in ICAM simulation = 400.

these arguments, drawing examples on water, fisheries and power management in the US Pacific North West. The Board on Sustainable Development [31] of the US National Research Council has advanced similar arguments. In the area of climate assessment the report of Working Group III of the IPCC [32] notes that in addition to our own work, the work of Bankes [33], Lempert et al. [34, 35], Laitner and Hogan [36], Van Asselt and Rotmans [37], and Yohe [38] has explored aspects of adaptive approaches to management.

In a paper in the journal *Risk Analysis* [3], we explore and illustrate the issue of dealing with models which one believes are only appropriate within a specified domain of the model parameter space. We also explore problems that can arise from the fact that when a complex model must be operated into a region of its phase space for which it was not designed, different elements of the model may degrade at different rates. For example, different elements of a time-stepped model may become unreliable at different times as the model is run far into the future. One strategy is to assess the probability of model failure as a function of time, or of some endogenous model variable (incremental radiative forcing in the case illustrated here). Then one can display a time series that reports the likelihood of model failure along with the time series of model output. Alternatively it may be possible to identify and disaggregate the various sources of model failure and treat them separately. The one serious problem with this strategy is that in complex real-world situations, experts are likely to be able to identify only a portion of all the limitations or "surprises" that could be encountered, and are likely to be willing to assess probabilities for only a subset of the total. Still, identifying some and getting part way to a full treatment is clearly better than simply ignoring the possibilities.

When it is known that one portion of a model will become unreliable more rapidly than other portions of the model (e.g., over time, the socioeconomic sub-model of ICAM will become unreliable before the geophysical model), it may be possible to develop much simpler order of magnitude models, or perform bounding calculations, which allow one to say something, even when detailed prediction is not possible. Through the use of subjective judgments, the results of several such analyses can be weighted and combined, in this case over time, to yield a more meaningful projection than would be obtained by running a detailed high-resolution model well past its domain of applicability. Before drawing conclusions from such results it would be wise to perform analysis of their sensitivity to the choice of subjective weights.

Figure 7 illustrates the general strategy we proposed. One starts with a detailed model that is likely to only be reliable for a few years. Gradually one moves over to a much simpler model based on order of magnitude considerations. Finally, in the long-term, one can only bound the result, without giving best estimates. The weighting functions that are used to combine models, and make the switch from one model to
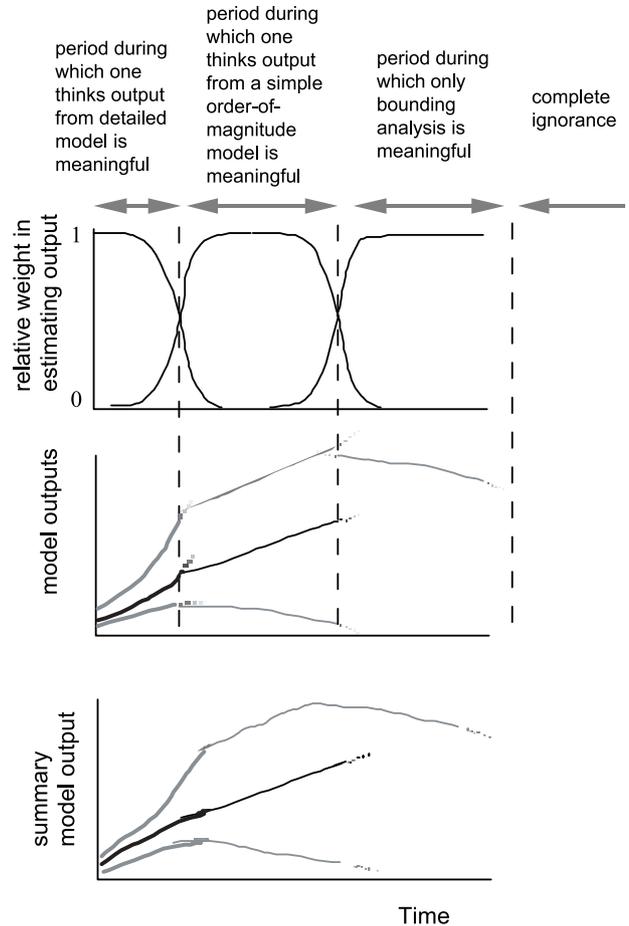

Fig. 7. Schematic illustration of the strategy of switching to progressively simpler models as one moves into less well understood regions of the problem phase space, in this case, over time. One starts with a detailed model that is likely to only be reliable for a few years. Gradually one moves over to a much simpler model based on order of magnitude considerations. Finally, in the long-term, one can only bound the result, without giving best estimates. The weighting functions for combining the models are based on subjective judgment. While the illustration shows three models over time, there is no reason why the number cannot be more or less than three. From [3].

another over time, must, of course, be based on subjective judgment. While the illustration shows three models over time, there is no reason why the number could not be more or less than three. Interested readers can find details and a specific application in [3].

5. CONCLUSIONS

I conclude by returning to the four questions which I posed at the beginning of Section 3 and offering my personal conclusions.

Question 1: How can we use uncertainty analysis, bounding analysis, and similar methods, to limit the range of possible "answers" even if perfect understanding and predictive capability cannot be achieved over the next few decades?

Answer: Investigators in this field are well on their way to doing this. As we move forward, we need to pay particular attention to:

– the fact that the ensemble of results from GCMs probably does not span the full range of uncertainty.
– the possibility that the future may not be simple linear extrapolations from the past.
– the fact that there is no ''objective'' way to state many relevant uncertainties, they are inherently subjective.

Question 2: How can we develop a clearer indication of what we can expect to know, and what we will probably not be able to know, about climate change and its likely impacts on the time-scale of the next few decades?

Answer: Doing this using expert subjective judgment poses no fundamental methodological challenge, but there is a very serious motivational bias. Investigators who have been enjoying substantial financial support by selling their research as promising answers, display reluctance to be too explicit about the likely limits of those answers since that might contribute to a reduction in the amount of future research funding.

Question 3: How can we use scientific knowledge and policy analysis to develop strategies for social and individual behavior and choice that are likely to be robust in an inherently uncertain world?

Answer: We have begun, but we need to do much more. We should abandon the (economists') fiction that we can find ''optimal policies'' and look instead for feasible policies. This includes:

– Acknowledging that consideration of equity are likely to be as or more important than efficiency.
– Looking for new and different ways to move forward (see, for example, [2]).

Question 4: As we elaborate our characterization of uncertainty, and what we cannot expect to know, how can we do it in such a way as not to provide aid to those who argue that no action is warranted until all uncertainties have been resolved?

Answer: We probably can't. However, the general public is pretty smart. Both they, and many corporate leaders, are coming around on this issue, even in the US! The public deal with uncertainty all the time. They have personal experience with making decisions in the face of uncertainty. When it is properly explained, they can understand both the basic uncertain nature of the climate problem and the need to take action in the face of that uncertainty.

ACKNOWLEDGMENTS

REFERENCES

1. Morgan, M.G., Pitelka, L.F. and Shevliakova, E.: Estimates of Climate Change Impacts on Forest Ecosystems. *Clim. Change* 49(3) (2001), pp. 279–307.
2. Morgan, M.G.: Managing Carbon from the Bottom Up. *Science* 289 (2000), p. 2285.
3. Casman, E.A., Morgan, M.G. and Dowlatabadi, H.: Mixed Levels of Uncertainty in Complex Policy Models. *Risk Anal.* 19(1) (1999), pp. 33–42.
4. Morgan, M.G., Kandlikar, M., Risbey, J. and Dowlatabadi, H.: Why Conventional Tools for Policy Analysis are Often Inadequate for Problems of Global Change. *Clim. Change* 41 (1999), pp. 271–281.
5. Morgan, M.G.: Uncertainty Analysis in Risk Assessment. *Human Ecol. Risk Assess.* 4(1) (1998), pp. 25–39.
6. Morgan, M.G. and Dowlatabadi, H.: Learning from Integrated Assessment of Climate Change. *Clim. Change* 34 (1996), pp. 337–368.
7. Morgan, M.G. and Keith, D.: Subjective Judgments by Climate Experts. *Environ. Sci. Technol.* 29(10) (1995), pp. 468–476.
8. Raiffa, H. and Schlaifer, R.: *Applied Statistical Decision Theory*, Harvard, 1961.
9. Howard, R.A. and Matheson, J.E.: *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA, 1983.
10. U.S. Nuclear Regulatory Commission.: *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants.* NUREG-75/014 (WASH-1400), Washington, DC, 1975.
11. Keeney, R.L.: Decision Analysis: An Overview. *Oper. Res.* 30 (1982), pp. 803–837.
12. von Winterfeldt, D. and Edwards, W.: *Decision Analysis and Behavioral Research.* Cambridge University Press, Cambridge, 1986.
13. Morgan, M.G. and Henrion, M.: *Uncertainty: A Guide to Dealing With Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge University Press, Cambridge, 1990.
14. Evans, J.S., Graham, J.D., Gray, G.M. and Sielken, R.L., Jr.: A Distributional Approach to Characterizing Low-Dose Cancer Risk. *Risk Anal.* 14 (1994), pp. 25–34.
15. Paté-Cornell, M.E.: Uncertainties in Risk Analysis: Six Levels of Treatment. *Reliab. Eng. Syst. Safety* 54 (1996), pp. 95–111.
16. Callahan, B.G. (ed.): Special Issue on Monte Carlo Analysis. *Human Ecol. Risk Assess.* 2 (1996), 1038 pp.
17. Presidential/Congressional Commission on Risk Assessment and Risk Management (1997), *Volume 1:Framework for environmental health risk management; Volume 2:Risk assessment and risk management in regulatory decision-making* (529 14th Street, NW, Suite 420, Washington, DC 20045).
18. Wallsten, T.S., Budescu, D.V., Rapoport, A., Zwick, R. and Forsyth, B.: Measuring the Vague Meanings of Probability Terms. *J. Exp. Psychol. General.* 155(4) (1986), pp. 348–365.
19. U.S. Environmental Protection Agency: *Proposed Guidelines for Cancer Risk Assessment.* EPA/600P-92/003C, Office of Research and Development, Environmental Protection Agency, Washington, DC, 1996.
20. Moss, R. and Schneider, S.H.: Uncertainties – Guidance Papers on the Cross Cutting Issues of the Third Assessment Report of the IPCC. In: R. Pachauri, R. Taniguchi, K. Tanaka (eds.): *World Meteorological Organisation*, Geneva, Switzerland, 2000.

21. National Assessment Synthesis Team: *Climate Change Impacts on the United States*. Cambridge University Press, Cambridge, 2000, 154 pp.

22. Budnitz, R.J., Apostolakis, G., Boore, D.M., Cluff, L.S., Coppersmith, K.J., Cornell, C.A. and Morris, P.A.: *Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and the Use of Experts*. Lawrence Livermore National Laboratory, URCL-ID-122160, Livermore, CA, 1995.

23. Kahneman, D., Slovic, P. and Tversky, A. (eds.): *Judgment Under Uncertainty*. Cambridge, 1982, 544 pp.

24. Hastie, R. and Dawes, R.: *Rational Choice in an Uncertain World*. 2nd ed. Sage Publications Inc., 2001, 372 pp.

25. National Defense University.: *Climate Change to the Year 2000: A Survey of Expert Opinion*, report published by the National Defense University in cooperation with the US Department of Agriculture, the Defense Advanced Research Projects Agency, the National Oceanic and Atmospheric Administration, and Institute for the Future, Washington, DC, February, 1978.

26. Stewart, T.R. and Glantz, M.H.: Expert Judgment and Climate Forecasting: A Methodological Critique of Climate Change to the year 2000. *Clim. Change* 7 (1985), pp. 159–183.

27. Nordhaus, W.D.: Expert Opinion on Climate Change. *Am. Scientist* 82 (1994), pp. 45–51.

28. Dowlatabadi H.: Sensitivity of Climate Change Mitigation Estimates to Assumptions About Technical Change. *Energy Econ.* 20 (1998), pp. 473–483.

29. Clark, W.: Witches, Floods, and Wonder Drugs: Historical Perspectives in Risk Management. In: R.C. Schwing, W.A. Albers (eds.): *Societal Risk Assessment: How Safe is Safe Enough*, Plenum Press, New York, 1980, 363 pp.

30. Lee, K.: *Compass and Gyroscope: Integrating Science and Politics for the Environment*. Island Press, 1993, 243 pp.

31. Board on Sustainable Development: *Our Common Journey: A Transition Toward Sustainability*. US National Academy Press, (1999), 363 pp.

32. IPCC: *Climate Change 2001: Mitigation*. Cambridge, 2001, 752 pp.

33. Bankes, S.C.: Exploratory Modeling for Policy Analysis. *Oper. Res.* 41(3) (1993), pp 435–449.

34. Lempert, R.J., Schlesinger, M.E. and Bankes, S.C.: When We Don't Know the Costs or the Benefits: Adaptive Strategies for Abating Climate Change. *Clim. Change* 33(2) (1996), pp. 235–274.

35. Lempert, R.J. and Schlesinger, M.E.: Robust Strategies for Abating Climate Change. *Climatic Change* 45(3/4) (2000), pp. 387–401.

36. Laitner, J.A. and Hogan, K.: Solving for Multiple Objectives: The Use of the Goal Programming Model to Evaluate Energy Policy Options. In: *Proceedings of the ACEEE Energy Efficiency in Buildings Summer Study*. American Council for an Energy Efficient Economy, Washington, DC, August 2000.

37. Van Asselt, M. and Rotmans, J.: Uncertainties in Perspective. In: J. Rotmans, B. de Vries (eds.): *Perspectives on Global Change*. Cambridge University Press, Cambridge, 1997, 479 pp.

38. Yohe, G. and Wallace, R.: Near Term Mitigation Policy for Global Change Under Uncertainty: Minimizing the Expected Cost of Meeting Unknown Concentration Thresholds. *Environ. Model. Assess.* 1(1/2) (1996), pp. 47–57.