Approaches for Performing Uncertainty Analysis in Large-scale Energy/Economic Policy Models

Antje Kann and John P. Weyant
Energy Modeling Forum, Stanford University
Terman Engineering Building, Room 406
Stanford, CA 94305

EMF Working Paper 14.6, June 1999
(appeared in Environmental Modeling and Assessment, Vol. 5, No.1, 2000, pp. 29-46)

Abstract

A number of key policy insights have emerged from the application of large-scale economic/energy models, such as integrated assessment models for climate change. These insights have been particularly powerful in those instances when they are shared by all or most of the existing models. On the other hand, some results and policy recommendations obtained from integrated assessment models vary widely from model to model. This can limit their usability for policy analysis. The differences between model results are mostly due to different underlying assumptions about exogenous processes, about endogenous processes and the dynamics among them, differences in value judgments, and different approaches for simplifying model structure for computational purposes. Uncertainty analyses should be performed for the dual purpose of clarifying the uncertainties inherent in model results and improving decision making under uncertainty. This paper develops a unifying framework for comparing the different types of uncertainty analyses through their objective functions, categorizes types of uncertainty analyses that can be performed on large models, and compares different approaches to uncertainty analysis by explaining underlying assumptions, suitability for different model types, and advantages and disadvantages. The appendix presents a summary of integrated assessment models for climate change that explicitly account for uncertainty.

Introduction

With rapid advances in computing power over the last decade, large-scale models have become essential to decision making in public policy. For example, policy makers in the area of global climate change rely on a variety of integrated assessment models[1] as an important tool for analyzing the consequences of different policies. These models, which originate from a group of globally distributed, diverse and interdisciplinary researchers, aide policy makers in determining actions to mitigate climate change. However, in some cases results and conclusions vary widely across the different models, which has made it difficult for scientists and policy makers to know how to use the results of integrated assessment models. Thus, policy makers' high demand for reliable model results has not yet been fulfilled. Other areas of policy analysis, such as health policy, are affected by similar issues of model inconsistency.

The disparities among model results can be ascribed to a few dominant factors:

- Different underlying assumptions about processes exogenous to the model. This is reflected in discrepancies among input parameters across the models.
- Different underlying assumptions about endogenous processes and the dynamics among them, resulting in different model structures.
- Differences in value judgments, such as the value of a human life.
- Different approaches for simplifying model structure for computational purposes.

In an effort to increase the use and usefulness of integrated assessment models in policy analysis, assumptions inherent in the different models need to be laid out explicitly, so that models can either be better synchronized or their differences be better understood. In addition to clarifying the assumptions on model inputs, modelers should be more explicit about their level of confidence in model outputs. Ideally, policy makers should be provided with recommendations for policies that are robust in the face of large uncertainty about future outcomes, and with suggestions on how to reduce that uncertainty efficiently.

Many modelers have spent considerable effort on fine tuning their models to replicate dynamics in physical and social sciences at levels that reflect current scientific knowledge and that can replicate historical data. The fact that both historical data and scientific knowledge of many dynamics are not exact has often been overlooked in the process of model improvement. Therefore, many observers have requested that the next stage in the model development process include a collective effort to better understand the assumptions and data that provide input to the models and to better describe the uncertainty inherent in the model outputs, with the ultimate goal of making model results more accessible to policy makers.

[1] Integrated assessment models combine scientific and economic models to allow for an integrated analysis of complex policy problems [IPCC 1995].

Understanding uncertainty is also the prerequisite to resolving uncertainty. Many types of uncertainty can be at least partially resolved by committing resources towards research efforts that represent uncertainty. In light of limited resources, policy analysts need to determine which research efforts represent the most effective and efficient use of resources.[2] Subjecting existing models to uncertainty analysis can help to determine which types of uncertainty should be addressed first.

What Defines an Uncertainty Analysis?

While our understanding of many physical and socioeconomic dynamics has progressed notably over the last few decades, new scientific areas that are not yet completely understood, such as HIV research and global climate change, have emerged. For example, there is still significant uncertainty about the nature of climate change, damages from climate change, and the costs of preventing climate change. Despite such large uncertainties, most integrated assessment models for climate change are deterministic, i.e., they perform one evaluation of a given state of the world and assumptions, reflected in a single set of input variables, thus ignoring uncertainty. Policy makers, on the other hand, need a measure of the robustness of the model outputs and conclusions to variations in model inputs. In addition, this information is useful only if it is presented in a manner comprehensible to policy makers.

Investigating the reliability of model results requires several steps, including:
- quantifying the extent to which output variables depend on variations in input,
- determining which of the input variables have the most significant effect on variations in outputs, and
- determining which inputs are least understood or least predictable.

Ideal results of an uncertainty analysis would include
- probability weighted values of the output variables,
- optimal decisions in light of imperfect knowledge,
- a measure of risk or dispersion about the outcome, and
- the value of information for key variables.

Uncertainty in models can be characterized into the following two general categories:
- parametric uncertainty, which arises due to imperfect knowledge, and
- stochasticity, which is due to natural variability in certain processes.

---

[2] For example, some advocates of integrated assessment models in policy analysis have proposed identifying a "portfolio of actions" to minimize the collective effects of climate change [IPCC, 1996]. Such a portfolio might include adaptation measures, emissions reduction, research and development on better technologies, and continued scientific research on the likelihood and effects of climate change. The latter activity is effectively the effort of reducing uncertainty. While some of the other actions might seem more tangible, reducing uncertainty through scientific research is a key contributor to increasing the effectiveness of the other measures.

As our knowledge about dynamics in the physical and social sciences improves over time, we expect to reduce parametric uncertainty.[3] On the other hand, natural variability will always occur, and stochastic uncertainty is not reduced over time.[4] Stochastic uncertainty can have a cumulative effect on the overall model uncertainty in problems with a long time horizon. In such cases the effect of stochastic variability may contribute more to outcome uncertainty than parametric uncertainty [Zapert et al., 1998].

There are some additional categories of uncertainty analysis, including uncertainty about values and uncertainty about model structure. A prime example of uncertainty about values in economic models are the value of a human life and the intertemporal discount rate. Policy choices that affect future generations tend to be very sensitive to the choice of discount rate, and most climate change models can obtain results on all ends of the spectrum by varying the discount rate. There is widespread disagreement about the "true" value of a discount rate and the extent to which it is tied to the rate of return on capital [Portney et al., 1999]. While the rate of return on capital is a stochastic uncertainty, the discount rate's link to the rate of return is a value uncertainty. Ideally, policies should be robust over a relatively large range of discount rates and other “value uncertainties”.

Uncertainty about model structure implies that a physical or socioeconomic process cannot be replicated accurately in a model, for reasons such as limited computing resources or incomplete understanding due to lack of empirical data. For example, there are various population growth models and different types of damage functions among which modelers can choose. Peck and Teisberg show that the form of the damage function (in addition to the choice of damage parameters) can greatly affect the optimal carbon control policy in their model: if damage is linear in temperature, little control is optimal, but if damage is a cubic power function of temperature rise, a high level of control is optimal by the end of the next century [Peck and Teisberg, 1992]. This result not only underscores the importance of identifying nonlinear responses in human and natural systems, but also in identifying the sensitivity of outcomes to different assumptions and submodels.

In some cases, modelers differ on what defines an uncertain parameter[5]. If it is possible to exert some control over an uncertain parameter through commitment of resources, there might be some contexts where such a parameter could also be considered a decision variable. For example, population growth and energy efficiency are two important uncertainties in climate models that can be influenced significantly. Spending funds on population control could have a

---

[3] Though sometimes we realize that we know even less about a parameter than previously believed. In such an instance, the measure of uncertainty about the parameter was inadequate and needs to be increased with the additional knowledge rather than decreased. See the section on Alternative Approaches: Higher Order Uncertainties for more on this subject.

[4] Note, however, that stochastic uncertainty is described in models through parameters, the knowledge of which may improve over time. Thus, even if variability remains constant over time, we may be able to improve our understanding of it and “reduce” some of the stochastic uncertainty in that way.

[5] See also van Asselt (1999) for a detailed treatment on different types and sources of uncertainty.

great impact on population growth in terms of lowering both the mean and the variance of that parameter, and spending funds on improving energy efficiency can result in more efficient technologies. Thus, rather than treating energy efficiency or population growth as uncertain parameters, the amount of money to spend on these factors could also be treated as a decision variable.[6]

Some Practical Implications

Ideally, all policy models should be subjected to an uncertainty analysis. However, depending on the type of model and analysis performed, uncertainty analyses can be complex and computationally intensive. As a result, many modelers have performed only very basic types of uncertainty analysis. To date, the major obstacle to performing all-inclusive uncertainty analyses is the limitation of computing resources, especially since many large models already face an important trade-off between the level of detail of analysis and run-time. A full uncertainty analysis may require hundreds of thousands of model runs, a task that is feasible for only the simplest of models. Thus, weaker model dynamics and less model detail could allow for more exhaustive uncertainty analysis [Parson and Fisher-Vanden, 1997].

In addition to the size of a model (such as the number of variables and complexity of internal dynamics), the model type also affects the feasibility of different kinds of uncertainty analysis. Most large-scale policy models, especially in the energy/economic field, can be classified either as *policy evaluation models* or *optimization models*. Policy evaluation models evaluate given policy scenarios and tend to be rich in physical detail, while optimization models optimize over key decision variables to achieve a certain objective, such as cost minimization or welfare maximization [IPCC, 1996]. Most policy evaluation models lend themselves more easily to uncertainty analysis than optimization models, since they require less computation than optimization models.

Other difficult issues that need to be considered when modeling uncertainty are:

- Model calibration: Most models are not calibrated for extreme cases of input parameters. Thus, they are implicitly designed to provide the most accurate answers when parametric uncertainties are relatively small.[7]
- Subjectivity: Descriptions of uncertainty can be somewhat arbitrary and may vary drastically across different experts [Tversky et al., 1974]. Some modelers provide “intervals of uncertainty” for important parameters, which represent ranges of values considered as appropriate by scientists but have no statistical interpretation [Braddock et al., 1995]. At the same time, parameters that describe uncertainty, such as high/low percentiles,

---

[6]However, it would take unrealistically large amounts of funding to convert these two parameters to true decision variables.

[7]For example, quadratic functions are generally good approximations for small excursions around point estimates, but they can be very poor approximations for large excursions, particularly if the underlying function exhibits non-linear or abrupt changes.

distributions, and discrete probabilities, strongly influence the outcome uncertainty[8]. In order to maintain comparability of policy recommendations and outcome uncertainties across models, consistent input uncertainties should be applied.

The remainder of this paper first presents a unifying framework that clarifies the relationship between the different types of models that incorporate uncertainty in mathematical terms on the basis of their objective functions. It then offers a guide on how to convert an existing deterministic model to a model that explicitly incorporates uncertainty. This guide describes different types of uncertainty analyses in detail and outlines underlying assumptions, advantages and limitations for each type.

FRAMEWORK FOR COMPARING MODELS WHICH INCORPORATE UNCERTAINTY ANALYSIS

We compare the different types of policy models that explicitly account for uncertainty by describing the relationship between their objective functions. The framework developed in this paper uses stochastic dynamic optimization (as explained later) as a unifying concept for analyzing uncertainty in large-scale models and then describes different variations of this general concept that are performed out of the practical needs described above. Thus, the framework proceeds from the most general and computationally intensive approach, to more restrictive and less complex approaches.

Stochastic dynamic optimization is an appropriate tool to model policy problems that are characterized by the interaction between actions taken today and consequences experienced tomorrow in the light of uncertainty. While stochastic dynamic programming theoretically represents the most comprehensive approach of analyzing uncertainty in this context, it usually does not pass the test of practicality. This motivated the approach of framing the other types of uncertainty analysis as special cases of this general concept that arise out of practical needs, each incorporating a different set of assumptions that makes computation feasible. These simplifications can take the shape of reducing model detail, restricting how uncertainty is modeled, or restricting how optimal choices are made. Figure 2 presents an overview of the framework:

[8] This paper is based on the assumption that probability is the most effective way to express uncertainty [Lindley, 1982].

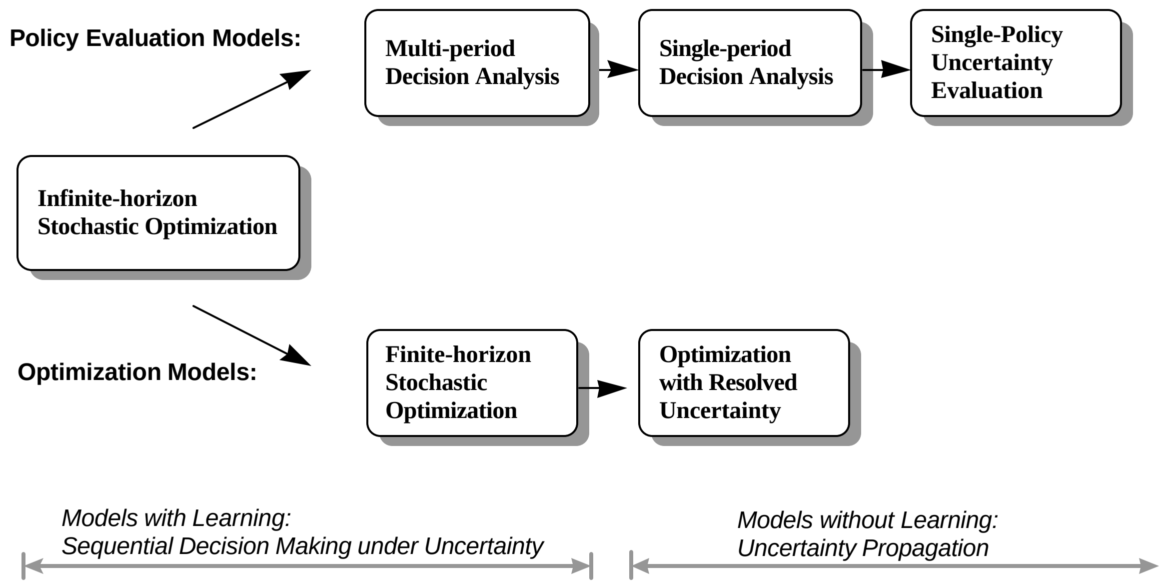

Figure 2: Complexity Hierarchy of Types of Uncertainty Analysis

The framework developed in this section is based on the premise that most policy models evaluate variables (some or all of which may be uncertain) over time, and then combine these variables into an overall performance measure which is maximized or minimized, or simply evaluated. This performance measure can vary widely from model to model. For example, most economic models sum the discounted utility of consumption over time, but different variables could also be added or averaged directly. In all cases, we shall refer to the evaluation of the performance measure as the objective function.

Our framework describes the objective function of a dynamic optimization model with uncertainty in its most general form and then points out the differences that arise when simplifying assumptions need to be made. As described above, the type of the underlying deterministic model (policy evaluation model versus optimization model) determines how uncertainty can be incorporated into an analysis. This framework incorporates these differences by describing representative objective functions for the different model types. We acknowledge the fact that some models might not be described completely by this framework, and that there may be hybrid models which may fall into several categories.

The variables found in most policy models can be grouped into the following categories:

- A set of state variables, which describe the state of the key variables, such as economic or climate indicators.
- A set of control variables, which describe the policy, e.g., allocated resources.
- A set of information variables describing uncertainty, such as moments of a distribution. Information variables are only present in models that account for learning, i.e. updating of information. In these models, information variables are observed at certain time periods which are the basis for revised decisions.

Figure 3 presents the framework, which is explained in the subsequent paragraphs.

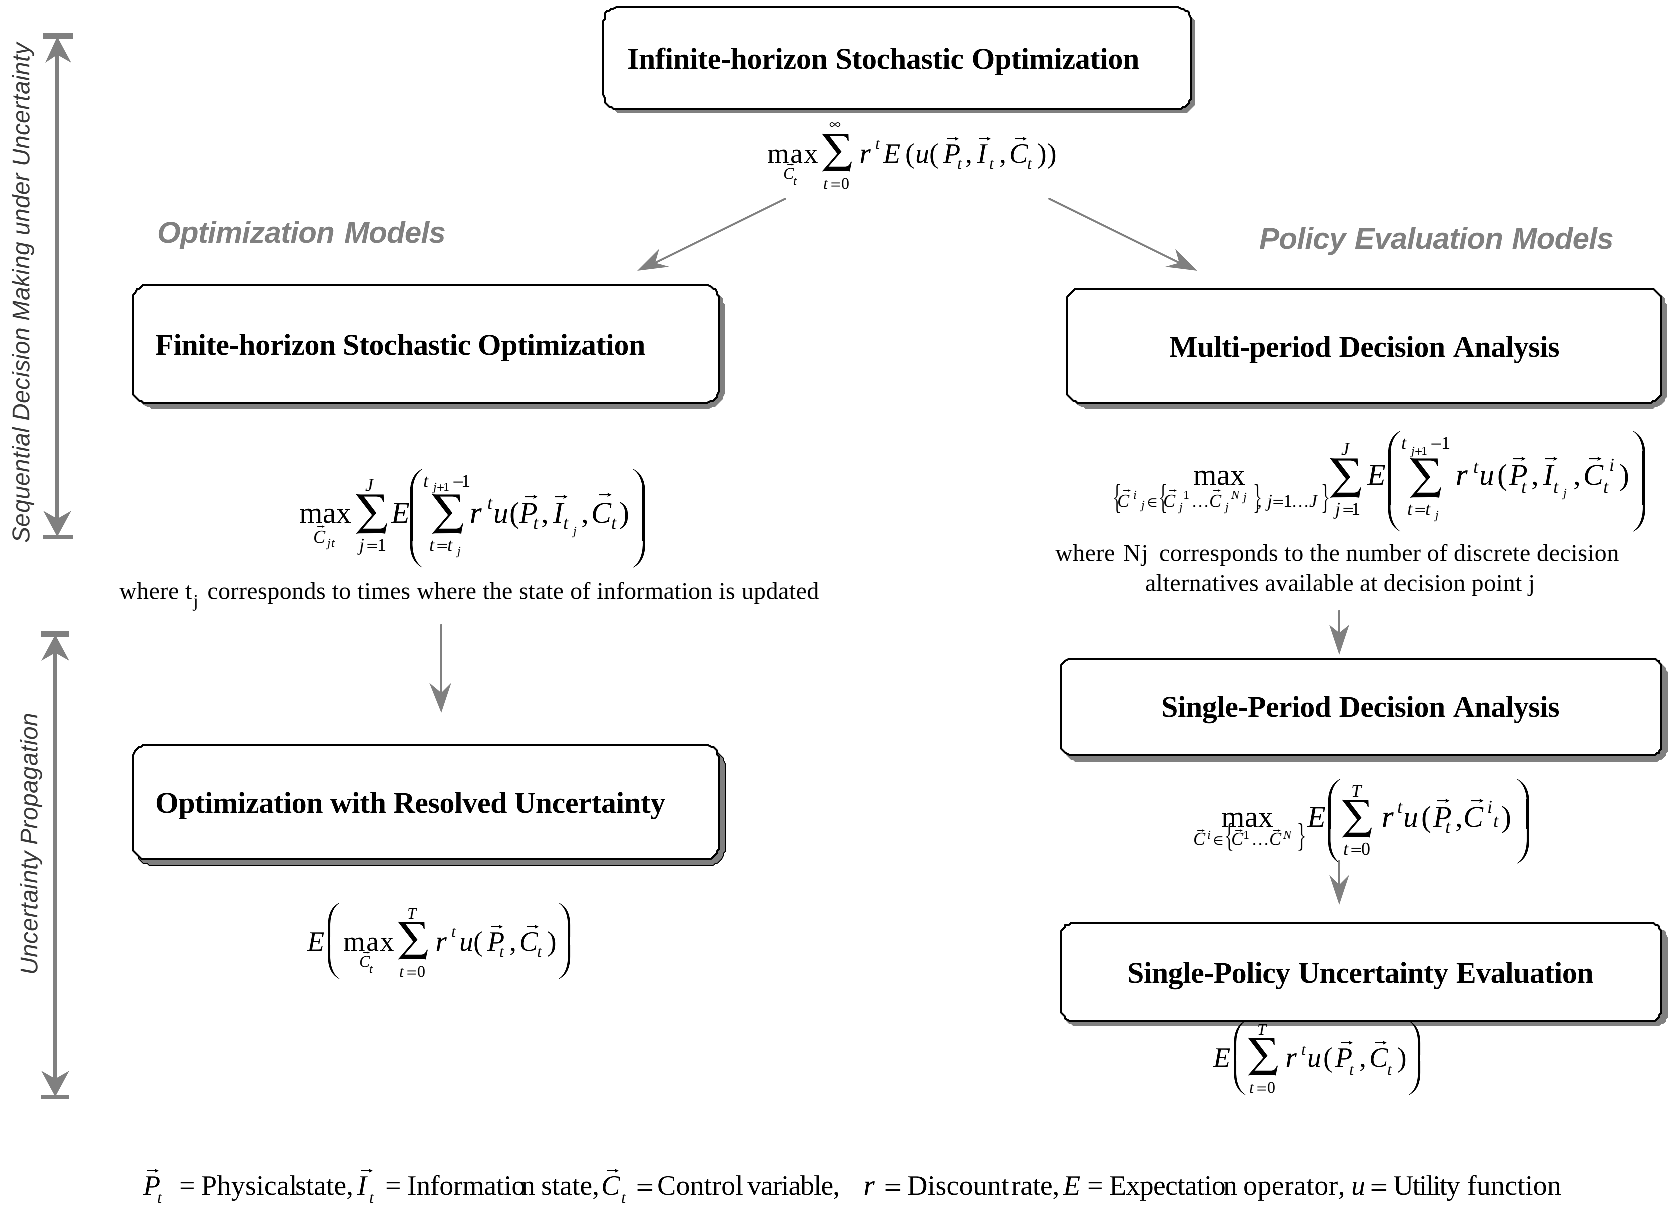

Figure 3: Framework for Analyzing Uncertainty

Sequential Decision Making under Uncertainty

Sequential decision making under uncertainty includes those models that determine optimal policies at more than one point in time. Models in this category range from a simple two-period decision analysis to an infinite-horizon stochastic optimization.

The most general objective function, which forms the basis for this framework, is that of stochastic optimization models which have an infinite horizon and allow for periodic decisions based on uncertain variables which are updated in each period:

$$\max_{\vec{C}_t} \sum_{t=0}^{\infty} r^t E(u(\vec{P}_t, \vec{I}_t, \vec{C}_t)) \qquad \text{(Infinite-horizon Stochastic Optimization)}$$

The decision maker maximizes the sum of the discounted expected utilities at each time period. The physical variables P and the information variables I are updated at every period t. Optimization of the control variables C_t takes place at every period, based on the information available at the time. The expectation at each time is based on the knowledge of all previous periods' variables, i.e., the physical variables and the updated information variables at $t-1$. This results in some very complex intertemporal dependencies.

The infinite time horizon implies that an analytic solution may be found, resulting in an optimal decision function based on the state of the vectors at each time [Stokey and Lukas, 1989]. As this solution process presents significant analytical challenges, a finite time horizon is often assumed. Finite-horizon problems can be solved numerically through dynamic programming. However, dynamic programming is computationally very expensive [Puterman, 1994], and the assumption that uncertainty is resolved at every period must be restricted as well.

While infinite-horizon stochastic optimization models can theoretically provide the most exhaustive and realistic analysis of policy problems, analytical and computational challenges generate the need for simpler, though more restrictive, approaches. In the following sections, we distinguish different approaches for analyzing uncertainty according to the nature of the underlying deterministic model (optimization model versus policy evaluation model).

Optimization Models

With the additional assumptions of a finite horizon and limited number of decision periods, the objective function for a *finite-horizon stochastic optimization* model can be represented as a special case of stochastic dynamic optimization:

$$\max_{\vec{C}_t} \sum_{j=1}^{J} E\left(\sum_{t=t_j}^{t_{j+1}-1} r^t u(\vec{P}_t, \vec{I}_{t_j}, \vec{C}_t) \right) \qquad \text{(Finite-horizon Stochastic Optimization)}$$

The vector of physical variables P is updated at every period, but the information variables and decisions are updated only at a limited number of predetermined periods denoted by {t_j,

$j=1\ldots J\}$. Note that the control variables may be different from period to period, but optimal policies are determined as sets $\{\vec{C}_{t_j} \ldots \vec{C}_{t_{j+1}-1}\}$ only at times t_j.

Policy Evaluation Models

The analogous case for policy evaluation models is referred to as *multi-period decision analysis*. The objective function for this type of problem is very similar to that of the stochastic finite horizon optimization model, with the exception of discrete (versus continuous) alternatives.

$$\max_{\left\{\vec{C}^i \in \left\{\vec{C}^1 \ldots \vec{C}^{N_j}\right\}, j=1\ldots J\right\}} \sum_{j=1}^{J} E\left(\sum_{t=t_j}^{t_{j+1}-1} r^t u(\vec{P}_t, \vec{I}_{t_j}, \vec{C}_t^i) \right) \qquad \text{(Multi-period Decision Analysis)}$$

Similar to stochastic finite horizon optimization models, the vector of physical variables P is updated at every period, and the information variables and decisions are updated at a limited number of J periods. The main difference is that the "optimization" takes place over a finite number N_j of discrete decision alternatives that are available at each of the J decision periods.

Uncertainty Propagation

When there is no learning involved, an optimal policy is determined only once, at the beginning of the model time frame. Models that incorporate uncertainty, but no learning, can be referred to as *uncertainty propagation* models. Again, a distinction is made between optimization models and policy evaluation models.

Optimization Models

The objective function of an optimization model which propagates uncertainty without learning can be represented as a special case of stochastic dynamic optimization:

$$E\left(\max_{\vec{C}_t} \sum_{t=0}^{T} r^t u(\vec{P}_t, \vec{C}_t) \right) \qquad \text{(Optimization with Resolved Uncertainty)}$$

In this case, optimization takes place at time 0, once for each possible uncertain state. Thus, the optimization is performed as if the true state of the world is known, and the expected value is taken over the results of these optimizations.

Policy Evaluation Models

The objective function of a policy evaluation model with discrete decision alternatives is similar, but evaluates the expectation and maximization operators in the opposite order:

$$\max_{\vec{C}^i \in \left\{\vec{C}^1 \ldots \vec{C}^N\right\}} E\left(\sum_{t=0}^{T} r^t u(\vec{P}_t, \vec{C}_t^i) \right) \qquad \text{(Single-period Decision Analysis)}$$

Since there is a finite number of discrete alternatives, the reduced computational complexity makes it possible to evaluate all possible states of the world for each decision alternative, i.e., the decision is made before uncertainty is resolved. This more closely resembles the situation of a decision maker at time zero. The "optimization" is performed to maximize the expected value of the outcome (as compared to taking the expected value of optimizations that were performed under the assumption that uncertainty has been resolved.)

Finally, the simplest type of uncertainty analysis consists of evaluating a given policy, without performing any optimization:

$$E\left(\sum_{t=0}^{T} r^t u(\vec{P}_t, \vec{C}_t)\right) \qquad \text{(Single-Policy Uncertainty Evaluation)}$$

In this case, the outcome of the model is the expected value of a given policy which is evaluated under uncertainty. Since this approach requires much less computation, it allows for the greatest detail in model resolution. Accordingly, the dimension of the vector of physical variables is typically much higher for models of this sort than for those where an optimization is performed.

The above framework describes different model types that incorporate uncertainty. To provide a practical context , the appendix contains a table that describes and categorizes integrated assessment models for climate change which explicitly account for uncertainty. Many of the models originated as deterministic models which were later adapted to include uncertainty analysis. For each model, the table lists the underlying deterministic model type, the type of uncertainty analysis performed, which uncertain variables are analyzed, and the main conclusions from the analysis.

The next section provides guidance on how to convert an existing deterministic model into a model that accounts for uncertainty.

CONVERSION OF DETERMINISTIC MODELS TO PROBABILISTIC MODELS

Preliminary steps

Depending on the size of the underlying deterministic model, it may be necessary to simplify the model structure for the purpose of performing an uncertainty analysis. For example, Braddock et al. simplify the very large and complex IMAGE model to a system of linear equations in order to obtain run-times that allow for a simulation of the model [Braddock et al, 1995]. The reduction in model detail and its corresponding impact on the accuracy of results needs to be carefully assessed.

Converting a deterministic model to a probabilistic model requires the selection of key uncertain parameters, and the robustness of the model to variations in these parameters needs to be understood. Sensitivity analysis and scenario analysis can provide useful insights for this purpose.

Sensitivity Analysis

Sensitivity analysis answers the question "how sensitive are model outputs to changes in model inputs?" It involves varying input parameters that are not known with certainty and recording the resulting changes in output variables and their effect on decisions. It can be used as a tool to identify which of the model parameters have the greatest effect on the output variables and results, to determine which parameters should be treated stochastically in further analysis, and to determine the break-even points between various alternatives for variations in a specific variable. Though it is not possible to model stochastic variability through this method, sensitivity analysis is covered here due to its essential role as a prerequisite to a more complex uncertainty analysis. A sensitivity analysis should be performed for every model before results are reported.

The simplest and most commonly performed analysis is the *single-value deterministic sensitivity analysis*. It involves setting each parameter of interest in a deterministic model to extreme points (usually its 5^{th} and 95^{th} percentile values) while holding all other variables at nominal values. The result is a measure of the functional relationship between a single parameter in question and output variables. For probabilistic models with non-random input parameters, *probabilistic sensitivity analysis* involves varying those parameters in the same manner. The output of such an analysis consists of a range of expected values and standard deviations.

When there are dependencies between variables, single-value sensitivity analysis does not sufficiently analyze uncertainty, as the overall uncertainty can be orders of magnitude larger than in the single variable case. *Joint sensitivity analysis* can produce a more accurate measure of output sensitivity by varying several parameters jointly. However, quantifying the interdependencies of variations in input variables can be a daunting task.

Sensitivity analysis is also useful for evaluating a model's sensitivity to those parameters which cannot or should not be modeled as uncertain. Specifically, sensitivity analysis is an appropriate tool for analyzing uncertainty about valuation and discounting, since it is less appropriate to specify probability distributions about these parameters.

Some of the shortcomings of sensitivity analysis include:

- The range of outcome values between the high and low percentiles might not reveal some of the uncertainty involved, especially if the maximum divergence from the best-guess value occurs in the interior of the range [Bankes, 1993].
- It is not possible to model stochastic variability through this methodology. Thus, it is not a substitute for performing uncertainty analysis.
- Performing a sensitivity analysis on a given model is based on the premise that the model structure is correct. It does not measure or detect specification error.[9]

[9] Most of the types of uncertainty analysis represented in this paper share this shortcoming.

Scenario analysis

Scenarios consist of combinations of different assumptions about possible states of the world, such as high population growth and low energy efficiency. Scenario analysis involves performing model runs for different combinations of assumptions and comparing the results. Scenario analyses can be classified as one of two types according to how the scenarios are generated: 1) scenarios are generated to look at interesting, meaningful, and varied combinations of states of the world [see Yohe, 1991], and 2) scenarios are generated from joint probability distributions on parameters according to technical criteria obtained from sampling methodologies [Iman and Conover, 1980; Nordhaus and Popp, 1996].

The first type of scenario analysis can provide very interesting insights and helps facilitate standardization and controlled comparison between assessments. It is especially useful when multiple models perform runs under identical input assumptions. This relieves different modeling projects of the need to choose and defend their own inputs [Parson, 1997]. However, as in the case of sensitivity analysis, this type of scenario analysis is not a substitute for uncertainty analysis. First, the scenarios are generally not weighted by probabilities. It is indeed difficult to determine the probability of a scenario, as one needs to determine the joint probability distribution of uncertain variables which tend to be characterized by high correlation. In addition, each model run is performed with best-guess values for a particular scenario, implying that the possibility of surprises is usually not captured.

In the second type of scenario analysis, scenarios are created by grouping uncertain parameters for computational purposes. This case will be covered in a later section.

Shortcomings of scenario analysis are:

- It is difficult to perform a rigorous uncertainty analysis using scenarios. Scenarios would need to be weighted probabilistically and would need to be mutually exclusive and exhaustive (i.e. include all uncertain states).
- Scenario analysis is still an area of active research. As a result, there is not much consensus as to the best scenario design, terminology often gets confused, and different types of analysis fall under this label. For example, the analysis of different policies is sometimes termed scenario analysis, while it is simply a policy evaluation exercise.

Once the model's sensitivity to variations in parameters is well understood, it can be converted to a probabilistic structure. While the mathematical framework presented in the previous section proceeds from the most general but complex to the more restrictive but simpler model types, the following section will apply a bottom-up approach, starting with the simplest kind of model structure and proceeding in the order of increasing complexity.

Propagating Uncertainty through a Deterministic Model

Perhaps the most commonly applied type of uncertainty analysis is the propagation of uncertainty through a deterministic model. The simplest implementation involves specifying a joint distribution (discrete or continuous) on a selection of input parameters and then propagating this uncertainty through to the model output. A more complex implementation involves modeling stochastic variability by reformulating a deterministic dynamical system into a stochastic dynamical system. This approach usually involves modeling certain variables as a stochastic process such as Brownian Motion. For an example, see [Zapert et al., 1998].

While a simple uncertainty propagation does not provide optimal decisions that consider uncertainty, the resulting distributions on output variables provide policy makers with a sense of the risk associated with the outcome. Such distributions are especially important when models are nonlinear. In nonlinear models the mean output value is not identical to the output value which corresponds to the mean of the input values. Similarly, propagating uncertainty is also important in the presence of risk aversion. Risk aversion implies that a decision maker values the average outcome at less than the average of the values for all outcomes. This could be the case when a decision maker is more interested in averting a particularly bad outcome under all circumstances than in selecting the policy which performs best on average.

In addition to a distribution on output variables, one can obtain measures of the relative importance of different input variables on the outcome, in the form of partial rank correlation coefficients or regression coefficients. Partial rank correlation coefficients are measures of the contribution of each uncertain input to the output uncertainty, after removing the effects attributable to other inputs [Iman and Conover, 1980]. After input uncertainties are propagated through a model and a distribution of the outcome is obtained, the partial rank correlation coefficients can be computed, from which input parameters can be ranked according to their effect on the outcome. An example of this type of analysis can be found in Hope et al. (1993), where partial rank order coefficients were obtained for a large number of uncertain input parameters to determine the important contributors to the uncertainty in the results, and the study identified strong regional differences in contribution of cost uncertainty to total uncertainty.

For computational purposes, propagation of uncertainty usually involves sampling from a joint distribution across input values. The most widespread technique used for this purpose is the Monte Carlo Method [Rubinstein, 1981]. However, for some models sampling from the full input distribution may be computationally too expensive, and it becomes necessary to sample from scenarios which summarize the distribution. The literature on sampling techniques includes suggestions for forming scenarios from the input distributions that will provide statistically accurate samples [Morgan and Henrion, 1990, Iman and Conover, 1980]. In this case, the combination of parameter values for scenarios does not necessarily represent an interesting state of the world, but rather a set of values that create the most statistically significant scenarios for sampling.[10] One of the methods used to compute scenarios is Latin Hypercube sampling [Iman and Conover, 1980]. This type of scenario analysis is sometimes performed in place of a full propagation of uncertainty (see [Nordhaus and Popp, 1996] for an example).

[10] However, there might be feasibility constraints on the combinations of parameter values selected in such a way.

To further decrease computing requirements, only a subset of variables could be treated stochastically. The selection of the subset should be based on an analysis of the importance of each variable, for example through rank-order correlation coefficients. It is also recommended to perform tests on what fraction of the overall uncertainty is represented by the subset of parameters [Nordhaus, 1994].

Other issues with propagating uncertainty are:

- It is difficult to specify joint distributions due to sometimes significant correlations between parameters. In the presence of strong interdependencies among variables, uncertainty could be grossly misrepresented if a separate distribution is specified for each variable.
- For computationally intensive models it might be impractical to specify probability distributions on all uncertain variables.
- Propagation of uncertainty yields different results for optimization models ("learn now then act") versus policy evaluation models ("act now then learn").
- Parameters can contribute to uncertainty but be irrelevant to decisions. For example, the outcome can vary greatly with changes in an input parameter, but all policy alternatives could vary in the exact same manner. A simple uncertainty propagation will not necessarily identify the policy-relevant parameters.[11]

Depending on whether the underlying deterministic model is a policy evaluation model or an optimization model, different methods and implications need to be considered.

Propagation of Uncertainty as Applied to Policy Evaluation Models

Single-policy propagation of uncertainty: The above general description of propagation of uncertainty most closely describes the analysis required for simple deterministic policy evaluation models. The output of this analysis is a measure of risk for the outcome of a deterministic model that evaluates a single policy, represented in a distribution on the possible states of the world, given today's uncertainty.

Single-period decision analysis: In addition to the above general description of uncertainty propagation, decision analysis formalizes the evaluation of different policies. A decision analysis evaluates a set of predetermined policies policy alternatives under consideration of probability distributions across uncertain variables. The evaluation can be based on multiple attributes of the decision criterion. The result is a recommendation of the optimal alternative (the alternative with the best expected value), given the current state of information [Raiffa, 1968].

Compared to continuous optimization models, decision analysis selects the optimal solution from a finite set of pre-specified alternatives. While this is a limitation, it also makes decision analysis an especially appropriate tool for analyzing those problems where a limited set of

[11] This may not be considered a drawback if stochastic variability propagation is not done to reveal policy-relevant parameters, but to assess the robustness of the model results.

decision alternatives is available. In situations where politics or other outside influences limit the policy choices, a decision analysis which evaluates the available alternatives can give more realistic recommendations than a model which optimizes over a set of continuous variables[12].

It is important to distinguish decision analysis from scenario analysis. Decision analysis incorporates some elements of scenario analysis, as a deterministic model is run for a variety of combinations of assumptions and decisions. However, a decision analysis incorporates probabilistic distributions across all uncertain states of the world and can analyze sequential decisions. Its output includes explicit distributions of the outcome variables under different policy alternatives, while sensitivity analysis is limited to ranking outcomes under a few scenarios of interest.

One problem inherent in decision analysis models is the "curse of dimensionality." For each uncertain variable, the dimension of the problem increases by a factor equal to the number of uncertain states. Depending on the number of uncertain variables and the number of their states, the underlying deterministic model might be run a very large number of times. The high dimensionality of most models also implies that each scenario cannot be individually tracked.

Other issues associated with decision analysis include:

- A decision analysis does not optimize over continuous variables and is therefore limited to evaluating predetermined policies. This lack of resolution on the decision alternatives can imply that the optimal policy recommended by a decision analysis might be a suboptimal solution.
- Obtaining joint distributions for events can be a very cumbersome process.
- The curse of dimensionality imposes limitations on the number of states and variables that can be incorporated into the model.
- Decision analysis cannot contribute to the selection of a suitable “portfolio of actions”, as desired by some policy makers, beyond comparing predetermined portfolios that are passed into the problem.

Propagation of Uncertainty as Applied to Optimization Models

Optimization with Resolved Uncertainty: Propagating uncertainty through optimization models is achieved with the same methodology as used for policy evaluation models, but represents a different type of analysis. Fundamental differences in the underlying deterministic models explain this difference: Policy evaluation models calculate an outcome for a given set of inputs, which implies that the result of a propagation of input uncertainties in a policy evaluation model is a distribution on the possible states of the world, given today's uncertainty. In contrast, optimization models search for the most efficient way to address a problem, given our current state of knowledge. The distribution that results from propagating uncertainty through an

[12] A continuous optimization model could theoretically achieve the same "reality bound" outcome by adding a set of constraints which restrict the outcomes to the same alternatives.

optimization model thus needs to be interpreted as follows: each point on the output distribution represents the result of an optimization for a particular uncertain state of the world represented in the input variables. This implies a "learn now, then act" approach in which the uncertain state is revealed before action is taken, which can be interpreted as optimization with resolved uncertainty.[13] In reality, the policy maker's situation is not one of resolving uncertainty now and acting optimally according to the revealed knowledge. Rather, the policy maker needs to optimize now in light of uncertainty, while knowledge is not revealed until a later time. A more appropriate approach of including uncertainty in optimization models would be to create a sequential decision making model.[14]

Adding Uncertainty to a Sequential Decision Making Model

The previously identified types of uncertainty analysis share one very strong assumption: The optimal policy is determined only once, today, with the currently available knowledge, and uncertainty is revealed after the decision has been made. A more realistic approach would take into account the fact that decisions are made nearly continuously over time as long-term uncertainty is reduced. In such a setting, a decision at each time need only be optimal for the period up to the next decision point. Since this more accurately describes the problem that most policy makers face, models should assist in identifying such short-term strategies in the face of long-term uncertainty, and analyze the effect of learning and adaptation on optimal policies .

Sequential decision making under uncertainty assumes that there are several points in time at which policy makers may make decisions that react to outcomes, and that their knowledge increases with time. At each such point in time, a decision is made based on a joint distribution that describes the possible outcomes that may occur during the following periods. Outcomes may be defined in this context either as states of the world or a new distribution (possibly with a different mean and/or with less spread than the previous distribution). In the subsequent period, another decision is made based on the updated knowledge, and so on. Thus, it is not necessary to assume that uncertainty is completely resolved at certain points in time, but rather that a probability distribution is updated. This process is frequently referred to in policy models as *learning*.

The Learning Process

Three main types of learning are represented in policy models [Kolstad, 1994]:

1) Active learning, whereby the effect of policy choices on certain key variables is observed for the purpose of obtaining information about uncertain parameters. For example, in the area of global warming policy, experiments such as perturbing emissions could reveal

[13] Knowing that policy *X* is the best policy in the state of the world *x* and that policy *Y* is the best policy in the state of the world *y* does not tell us how to choose an optimal policy before the state of the world is revealed [Ermoliev et al., 1988].

[14] However, if the optimal solution is insensitive to variations in the uncertain parameters, optimization with resolved uncertainty can indicate that the policy recommendation is robust in the face of uncertainty.

information about uncertain parameters by observing their effect on the economy and the climate system. [15]

2) Purchased learning, whereby knowledge is purchased (e.g., by allocating resources to R&D efforts).
3) Autonomous learning, where the mere passage of time reduces uncertainty.

Delaying action in order to learn more involves several tradeoffs:

- Waiting to learn more versus preventing irreversible damage: When learning is possible, and in the absence of irreversible outcomes, it tends to be advantageous to delay action until some uncertainty is reduced [Dixit and Pindyck, 1994]. On the other hand, the less learning is anticipated and the more irreversibilities are present, the more aggressive the short-term strategy should be.
- Waiting to learn more versus higher costs later: Even when damages are not irreversible, the mitigation costs may increase during the waiting period. In this case, the benefits from increased knowledge may be offset by the increased costs.
- Waiting to learn more versus beneficial effects from early action: Early policy action may have additional benefits such as triggering private sector R&D and innovation. Thus, the rate of uncertainty resolution could actually be increased by not waiting to learn more.

Issues surrounding learning abound in the field of climate change policy. The literature on integrated assessment models reveals that learning is usually modeled as an endogenous process (type 3), and that for computational purposes, most models have not considered more than 2 periods of learning. Some climate models have shown, and intuition supports the notion, that policy strategies which are adaptive and use results from learning to change the decision will on average perform significantly better than one-time optimal strategies which optimize over the entire time period of interest [Lempert et al., 1996]. Thus, it is usually preferable to adopt a sequential decision process and update decisions as new information is obtained. The decisions made before all uncertainty is resolved are suboptimal relative to sequential decisions made with perfect information [Kelly and Kolstad, 1996].

The sequential decision making framework also makes it possible to identify *hedging strategies*, which balance the risk of waiting with those of premature action [Manne and Richels, 1995]. Hedging can be viewed as building contingency plans and responding to opportunities and dangers as they become apparent, as opposed to averaging the different policies which are optimal for different states of the world. In the field of climate change, hedging strategies could be part of the portfolio of actions that minimize the risks of climate change, e.g., R&D programs for non-emitting technologies that would increase the likelihood that abatement costs would drop in the future [Lempert et al., 1996].[16]

---

[15] In reality, experiments cannot be conducted for many variables due to irreversibility, long time lags, and detection problems [Hammit, 1994].

[16] An Energy Modeling Forum study compared the application of hedging strategies across seven Integrated Assessment models. The study showed varying sensitivity of results to hedging strategies across the different models, but generally showed that some delay of action is optimal. However, the

Other issues associated with the learning process described above are:

- Due to computational issues, the process of updating information needs to be modeled at discrete points in time, which tends to greatly oversimplify the decision making process. In reality, decisions are not made every 20 or 50 years, but adjustments to policies are made continuously as information is updated. Whether or not a policy is revised may also depend on the state of the information revealed.
- The learning process is not straightforward, but rather obstructed by noise, natural variability, measurement errors and imperfect understanding of social and physical dynamics. An adaptive strategy will outperform one-time optimal strategies only if the observed phenomena will yield meaningful information [Lempert et al., 1996].

Expected Value of Information

One of the useful by-products of sequential decision making under uncertainty is the ability to identify the *expected value of information* for key uncertain variables. The value of information for an uncertain variable is determined by subtracting the model's expected net benefit when the uncertainty about the variable is not resolved until after the decision is made (or optimization is performed), from expected net benefits when uncertainty about the variable is resolved before the decision is made (or optimization is performed). In the former case, a policy decision is made based on the expected outcome, whereas in the latter case, different policies are chosen for each known state of the world [Howard, 1966].

The value of information is a useful tool in assessing the value of resolving the uncertainty about an uncertain variable. If it was possible to completely reveal the true value of an uncertain variable for a given price, then one should pay this price if it is lower than the value of information for that variable. For those models whose objective function measures the utility of consumption, the expected value of information can also be interpreted as the amount of consumption that society is willing to forego in exchange for information on the true state of the uncertain variables [Chao, 1995].

Depending on the nature of the uncertain variables on which information can be obtained, synergistic effects among the variables should be explored. When several uncertainties are resolved at once, their joint value of information may be very different from, and in certain instances much larger than, the sum of the individual values [Peck and Teisberg, 1993, Nordhaus and Popp, 1996]. If the joint value of information dominates the sum of the individual values, policy makers would gain higher benefit from performing comprehensive research rather than narrowly focused research.

In most cases, we do not know at what point in time uncertainties will be resolved. Another important question is then whether the resolution of uncertainty is urgent or whether strategies

---

analysis also showed that most cases call for some degree of immediate action. The optimal degree of immediate action was stronger than in the deterministic case [Manne, 1995].

are robust to the timing of uncertainty resolution. If the timing is urgent, policy makers should spend more funds improving knowledge through research, whereas if strategies are robust to timing, more funds can be allocated to other activities such as mitigation and adaptation [Nordhaus, 1994]. The *expected value of early revelation of uncertainty* is computed by determining the value of information at different times of resolution of uncertainty and subtracting it from the value obtained at time zero. The differences in these values reveal the value of early knowledge.

Computing the expected value of information raises the following issues:

- The value of information for a certain variable can exhibit large differences across different models. These differences represent either the emphasis that different models place on certain variables, or the dispersion of the distribution that was assigned to the variable. Uncertainty ranges are often subjective, but have large impact on the value of information.
- Computing the value of information individually for different variables might not provide answers to the questions that policy makers face. Most likely, policy makers will not trade off studying natural science versus economic impacts, but rather will need to allocate funds across the board. However, computing the joint value of information for several uncertain variables may require calculations that are too complex to perform, especially in the presence of correlations among the variables.
- The expected value of information is not necessarily correlated with the importance of a variable as determined by a deterministic sensitivity analysis. In other words, the variable to which utility is most sensitive is not the one with the highest value of information, if knowledge of that variable does not change the optimal policy much. New information is most valuable for parameters most closely related to policy making [Nordhaus and Popp, 1996].
- The information obtained is usually not perfect, which makes the concept of eliminating uncertainty unrealistic. To be more accurate, one should determine the value of reducing uncertainty from one level to another, less uncertain, level.
- In certain policy areas such as climate change the identity of the decision maker who would place value on the information is not clear.[17]

Sequential Decision Making under Uncertainty, as Applied to Policy Evaluation Models

The typical framework for sequential decision making under uncertainty, as applied to policy evaluation models, is *multi-period decision analysis*. Multiple sequential decisions can usually be analyzed with relative ease through this methodology. Decision analysis also provides a straightforward approach for solving for the expected value of information. All comments in the above section on single-period decision analysis apply to the multi-period setting as well.

---

[17] While this issue is raised frequently in connection with the value of information, it is actually a more general problem that underlies most modeling efforts (e.g., whose utility function, discount rate, etc. should be used).

Sequential Decision Making under Uncertainty, as Applied to Optimization Models

This category is often referred to as *stochastic optimization* and encompasses a variety of methodologies. Optimization models tend to be computationally intensive, even when uncertainty is not modeled explicitly. In the context of sequential decision making under uncertainty, the complexity increases by multiple dimensions, as optimizations have to be performed for every uncertain state of the world at every point in time where uncertainty is resolved. Accordingly, the many different approaches to stochastic optimization represent different tradeoffs regarding detail in the description of uncertainty, the type of optimization, the dimension of tractable variables, and run-time.

The main difficulties associated with this class of models are:

- While policy evaluation models can easily be converted to multi-stage decision analysis models[18], optimization models require structural modification to incorporate multi-stage uncertainty.
- Due to the computational complexity of optimization models, sequential decision making under uncertainty can only be performed for a very limited set of uncertain states of the world. This represents a tradeoff between resolution in policy and exhaustive representation of uncertainty.
- Due to computational complexity and long run times, certain optimization models may need to neglect various feedback loops among the different elements of physical and socioeconomic systems once uncertainty analysis is included.
- When a stochastic constraint is binding, the high shadow price associated with the constraint can affect the outcome disproportionately and result in unrealistic recommendations.

Solution methods depend strongly on the number of stages at which uncertainty is resolved and on the choice of time horizon of the problem, i.e. finite or infinite [Kloeden et al., 1993].

Two-stage Stochastic Optimization

The most common approach to sequential decision making under uncertainty, as applied to optimization models is a two-stage model, where the first stage consists of decisions taken before the uncertainty is resolved, and second stage decisions are taken after uncertainty has been resolved. The set of second stage decisions can be different depending on the outcome of the experiment. The class of models in which some decisions or recourse actions can be taken after uncertainty is disclosed is also referred to as *recourse programs.* An exhaustive treatment of two-stage stochastic linear programs with fixed recourse is given in Birge et al. (1998).

To perform such an analysis, the model is transformed into a set of parallel optimization problems. Each of the parallel problems represents a different state of the world with an appropriately adjusted constraint set. The outcome values of the parallel problems are then combined by a probability weighted objective function. An additional set of constraints is

---

[18] Essentially, decision analysis can be perceived as a "front end" that feeds parameters into the deterministic model.

necessary to ensure that the solution is the same for all states of the world during the first stage, i.e. the time frame before uncertainty is resolved.

In the field of climate change, several optimization models have been included in a collective effort of developing simple two-period optimization models, labeled as "act then learn" models. Most modeling groups have limited this type of analysis to two uncertain states of the world [Manne, 1995].

Multi-Stage Finite-horizon Stochastic Optimization

This approach is a generalization of the two-stage approach. However, due to the increased modeling complexity, it requires different solution techniques. At each stage, all possible states of the world need to be considered, and the decision variables depend on the realization of the stochastic variables in each previous period. Multi-stage optimization problems are typically solved by dynamic programming.

One of the solution methods employed in stochastic optimization is *dynamic programming*. Dynamic programming presents the solution for each decision period compactly as a pair consisting of the optimal action as a function of the current state of the world and the expected value of all future actions, assuming the optimal action is followed at each time [Howard, 1960; Ross, 1983]. With a finite time horizon, the problem can be solved recursively from the last time period backwards, though the curse of dimensionality limits the number of time periods and number of stochastic parameters for which this is feasible. With large state spaces, this approach becomes quite computationally cumbersome [Puterman, 1994], and typically, only vectors of very small dimensions can be accommodated.

An example of a deterministic linear program that was converted to a stochastic dynamic programming model can be found in Fragniere et al. (1995).

Infinite Horizon Stochastic Programming

Another approach is to consider an infinite horizon and use discounting to establish a stationary policy so that one need only find an optimal decision associated with a state for any stage. In the case of an infinite time horizon, the model can be constructed to find an analytical solution consisting of a decision function which is employed in each period and a value function. The decision at each stage depends only on the uncertain state of the world [Howard, 1960]. The solution space for infinite horizon stochastic dynamic programming models is a function space, which is infinite dimensional. Finding a solution function can be very difficult and requires some strong assumptions on the function space [Stokey and Lucas, 1989]. Usually, only one or two parameters can be modeled in such a way, given current computing resources [Kelly and Kolstad, 1996].

Most infinite-horizon stochastic optimization can be distinguished by the assumption on whether time is continuous or modeled in discrete stages. Discrete time models often use Markov processes to model the random movements of variables (see Stokey and Lucas, 1989), whereas continuous time models represent uncertain variables in continuous-time stochastic processes, such as Brownian Motion [Harrison, 1985].

Issues associated with infinite-horizon stochastic optimization are:

- The solution techniques for an infinite horizon model require stationarity of the problem, i.e., the parameters describing the stochastic process may not change over time.
- This process is analytically very challenging and requires mathematical expertise, thus limiting its accessibility to many modelers. In addition, solution functions may not exist for the problem to be analyzed.
- Determining parameters for a stochastic process may require statistical methods that are more difficult than those involved with defining probability distributions on uncertain variables.

Despite the computational intensity and the resulting limitations on the number of parameters that may be modeled stochastically, the infinite horizon approach offers one major advantage: Once a solution function has been found, further analyses can be performed without complicated computations. Since the solution is an analytical function of model inputs, the effects of variations in input values may easily be examined [Kelly and Kolstad, 1996]. Thus, stochastic optimization with an infinite horizon offers a way to incorporate stochasticity in key processes into the analysis, and at the same time offers a simple way to perform sensitivity analysis on other parameters.

Alternative approaches

Cultural theory

Cultural theory specifies perspective-based alternative model routes in which not only parameters but also relationships are varied according to the bias and preferences of a particular perspective. This results in alternative model structures [Van Asselt and Rotmans, 1995]. While cultural theory can provide fresh perspectives on the types of uncertainties, it does not yield quantitative results.

Exploratory modeling

Exploratory modeling extends uncertainty analysis beyond parametric variations. In traditional uncertainty analysis, input parameters are varied along a distribution. Exploratory modeling, however, emphasizes that the model itself should not be taken for granted and explores around different model variants. Thus, it can reveal regimes of qualitatively different behavior which result from excursions over non-linear functions and/or functional uncertainty about the system under study [Bankes, 1993].

The drawbacks of exploratory modeling are:

- It requires a high degree of analysis, and is thus impractical for large models.

- It does not provide a distribution of outcome variables because the sampling across the uncertainty space is non-random.

Higher order uncertainties

Not all distributions about uncertain parameters are equally well understood, some being more ambiguous than others. Ambiguity is a notion of exactness in the parameters of a probability distribution. The approaches to uncertainty analysis presented in this paper do not account for different types of uncertainties. Well understood uncertainty is considered no different than ambiguous and very subjective uncertainty.

Paul Fischbeck has proposed a theory to incorporate ambiguity, i.e., uncertainty about probabilities, into the normative framework of decision analysis, which also accounts for decision maker's preferences towards ambiguous probabilities. This method uses a multiple level probability framework which captures a complete uncertainty description of any event. By maximizing the multiple level expected utility, the decision maker can account for risk and ambiguity [Fischbeck, 1991].

The main drawback of this theory is that each level of uncertainty adds another dimension to the problem. In a situation where dimensionality is already the main problem, it would not be feasible for all modelers to perform this type of analysis.

Minimax Regret Strategies

This approach finds optimal strategies that take into account multiple states of the world, but which do not use expected utility as the criterion. As the only requirement is identifying the possible states of nature, it thus avoids the subjective task of estimating probabilities. The most common way to implement this approach, which is also known as Robust Programming is to analyze the regret of different policies. The regret of a policy is defined as the difference between the payoff of the optimal policy under perfect information and that with the given policy, and it is determined for each uncertain scenario. The policy which has the smallest regret under consideration of all uncertain states is optimal [Savage, 1951]. An application to a large energy policy model can be found in Kanudia et al. (1998).

Conclusion

This paper presented an overview of different types of uncertainty analysis that have been performed on integrated assessment models and introduced a framework for comparing the different types of uncertainty analyses through their objective functions. Due to the different assumptions and implications inherent in each model, there are limitations on what types of uncertainty analysis can be performed. The main purpose for developing a unifying framework was to clarify the different assumptions and limitations that underlie each type of analysis.

The main obstacle to performing a full uncertainty analysis which includes all variables is the lack of sufficient computational resources. Given current limitations, a tradeoff needs to be made among model detail and uncertainty analysis. If a problem is characterized by significant uncertainty or potentially important feedbacks, computer resources might be better spent on exploring a large number of alternative problem formulations than to increase the resolution in the best-estimate model.

It is conceivable that at some point in the future, stochastic dynamic optimization could be performed for integrated assessment models with many uncertain parameters. In the meantime, sequential decision making under uncertainty is one methodology that can provide reasonable comparisons across model types.

One goal of performing uncertainty analyses is to increase the usefulness of integrated assessment models to policy makers. This goal should be kept in mind when communicating the results of an uncertainty analysis. Quantifying the many types of uncertainty about a model and the underlying processes and values can be a daunting task, and communicating all the results can leave the end-user confused. While modelers should perform as many types of uncertainty analyses as their resources will sustain, only a subset of these should be considered for publication of results to policy makers, while others should be performed mainly as good modeling practice and to increase confidence in model structure and choice of parameters.

References:

Bankes, Steve C., Exploratory Modeling for Policy Analysis, Operations Research, Vol. 41, No 3, May-June 1993.

Birge, John R. and Francois Louveaux, Introduction to Stochastic Programming, Springer-Verlag, New York, 1997.

Braddock, R. D., J. A. Filar and R. Zapert, Stochasticity in the Image Greenhouse Model, Mathematical Computer Modeling, Vol. 22, No. 10-12, 1995.

Chao, Hung-Po, Managing the Risk of Global Climate Catastrophe: An Uncertainty Analysis, Risk Analysis, Vol. 15, No. 1, 1995.

Cohan, David, Stafford, R., Scheraga, J., Herrod, S., The Global Climate Policy Evaluation Framework, A&WMA, 4/1994.

Dixit, Avinash K. and Robert S. Pindyck, Investment under Uncertainty, Princeton University Press, Princeton, 1994.

Dowlatabadi, Hadi and Matt Ball, An overview of the Integrated Climate Assessment Model Version 2, (ICAM-2), presented at the Western Economics Association Conference, 6/29/94.

Ermoliev, Yu. and R. J-B Wets, Stochastic Programming: An Introduction, in Yu. Ermoliev and R. J-B Wets, editors, Numerical Techniques for Stochastic Optimization, Springer-Verlag, New York, 1988.

Filar, Jerzy A., and Radoslaw Zapert, Uncertainty Analysis of a Greenhouse Effect Model, in C. Carrare and A. Haurie, eds., Operations Research and Environmental Management, Kluwer Academic Press, Dordrecht, 1996.

Fischbeck, Paul S., Epistemic Uncertainty in Rational Decision Making, Ph.D. Dissertation, Stanford University, 1991.

Fragniere, E. and A. Haurie, MARKAL-Geneva: A Model to Assess Energy-Environment Choices for a Swiss Canton, in C. Carrare and A. Haurie eds., Operations Research and Environmental Management, Kluwer Academic Press, Dordrecht, 1996.

Grubb, Michael, Technologies, energy systems and the timing of CO2 emissions abatement, Energy Policy, Vol. 25, No. 2, 1997.

Hammit, James K., Outcome and Value Uncertainties in Global Change Policy, December 1994.

Harrison, J. Michael, Brownian Motion and Stochastic Flow Systems, Wiley, New York, 1985.

Hope, Chris, John Anderson and Paul Wenman, Policy analysis of the greenhouse effect- An application of the PAGE model, Energy Policy Vol. 21, No. 3, 1993.

Howard, Ronald A., Information Value Theory, IEEE Transactions on Systems Science and Cybernetics, Vol. 2, No. 1, August 1966.

Howard, Ronald A., Dynamic Programming and Markov Processes, MIT Press, Cambridge, MA,1960.

Iman, Ronald L. and Conover, W.J., Small Sample Sensitivity Analysis Techniques for Computer Models, with an Application to Risk Assessment, Communications in Statistics, A9(17), 1980.

Intergovernmental Panel on Climate Change (IPCC), Climate Change 1995, Economic and Social Dimensions of Climate Change, Cambridge University Press, 1996.

Kanudia, Amit, and Richard J. Loulou, Minimax Regret Strategies for Greenhouse Gas Abatement: Methodology & Application, presented at INFORMS Montreal, April, 1998.

Kelly, David L., and Charles D. Kolstad, The Climate Change Footprint: Will we see it before it is upon us? Draft, March 1996.

Kloeden, P. E., E Platen and H. Schurz, Numerical Solutions of SDE through computer experiments, Springer Verlag, New York, 1993.

Kolstad, Charles D., George Bush vs. Al Gore, Irreversibilities in greenhouse gas accumulation and emission control investment, Energy Policy, 1994.

Lempert, Robert J., Michael E. Schlesinger and Steve C. Bankes, When we don't know the costs or the benefits: Adaptive Strategies for Abating Climate Change, Climatic Change, 1996.

Lindley, Dennis V., Scoring Rules and the Inevitability of Probability, International Statistical Review, Vol. 50, 1982.

Manne, Alan S. and Timothy R. Olsen, Greenhouse Gas Abatement-toward Pareto-Optimal Decisions under Uncertainty, January 1994.

Manne, Alan and Richard Richels, The Greenhouse debate—Economic Efficiency, Burden Sharing and Hedging Strategies, April 1995.

Manne, Alan S., Stanford University, Hedging Strategies for Global Carbon Dioxide Abatement: A Summary of Poll Results, EMF 14 Subgroup—Analysis for Decision under Uncertainty, 12/95.

Morgan, M. Granger and Max Henrion, Uncertainty: A guide to dealing with Uncertainty in Quantitative Risk and Policy Analysis, Cambridge University Press, New York, 1990.

Nordhaus, William D., Managing the Global Commons, MIT Press, Cambridge, 1994.

Nordhaus William D., and David Popp, What is the Value of Scientific Knowledge? An application to Global Warming using the PRICE Model, 2/19/96.

Parson, Edward A. and Fisher-Vanden, Karen, Integrated Assessment of Global Climate Change, Annual Review of Energy and the Environment, Volume 22, 1997.

Peck, Stephen C. (EPRI) and Thomas J. Teisberg, CETA: A Model for Carbon Emissions Trajectory Assessment, The Energy Journal, Vol. 13, No. 1, 1992.

Peck, Stephen C. (EPRI) and Thomas J. Teisberg, Global Warming Uncertainties and the Value of Information: An Analysis using CETA, Resource and Energy Economics 15, 1993.

Peck, Stephen C (EPRI) and Thomas J. Teisberg, Optimal CO2 Control Policy with Stochastic Losses from Temperature Rise, Climatic Change 31: 19-34, 1995.

Plambeck, Erica L. and Chris Hope, PAGE 95, an updated valuation of the impacts of global warming, Energy Policy, Vol. 24, No. 9, 1996.

Portney, Paul R. and John P. Weyant, eds., Discounting and intergenerational equity, Resources for the Future, Washington, 1999.

Puterman, M. L., Markov Decision Processes: Discrete Stochastic Dynamic Programming, John Wiley, New York, 1994.

Raiffa, Howard, Decision Analysis: Introductory Lectures on Choices under Uncertainty, Addison-Wesley, Reading, MA, 1968.

Ross, Sheldon M., Introduction to Stochastic Dynamic Programming, Academic Press, New York, 1983.

Rubinstein, R. Y., Simulation and the Monte Carlo Method, John Wiley, New York, 1981.

Savage, Leonard J., The Foundations of Statistics, New York, Wiley, 1954.

Stokey, Nancy L. and Robert Lucas, Jr., Recursive Methods in Economic Dynamics, Harvard University Press, Cambridge, 1989.

Tol, Richard S. J., Tvd Burg, HMA Jansen, H. Verbruggen, The Climate Fund, Some notions on the socioeconomic impact of greenhouse gas emissions and emission reduction in an international context, Institute for Environmental Studies, RG5/03, Vrije Universiteit, Amsterdam, 1995.

Tol, Richard S. J., A Decision-Analytic Treatise of the Enhanced Greenhouse Effect, Institute for Environmental Studies, RG5/03, Vrije Universiteit, Amsterdam, 1997.

Tversky, Amos and Daniel Kahneman, Judgment under Uncertainty: Heuristics and Biases, Science, Vol. 185, September 1974.

van Asselt, Marjolein B. A. (1999). "Uncertainty in Decision Support: From problem to challenge." ICIS working paper I99-E006, International Centre for Integrative Studies (ICIS), University of Maastricht, the Netherlands.

van Asselt, Marjolein B. A. and Jan Rotmans, Uncertainty in Perspective, Global Environmental Change, Vol. 6 No. 2, 1996.

van Asselt, Marjolein B. A. and Jan Rotmans, Uncertainty in Integrated Assessment Modelling-a Cultural Perspective-Based Approach, 3/1995.

Yohe, Gary and Rodney Wallace, Near term mitigation policy for global change under uncertainty: Minimizing the expected cost of meeting unknown concentration thresholds, Environmental Modeling and Assessment, Vol. 1, No 1,2, June 1996.

Yohe, Gary W., Selecting “interesting” scenarios with which to analyze policy response to potential climate change, Climate Research, December 1991.

Zapert, R., Gaertner, P. S., Filar, J. A., Uncertainty Propagation within an Integrated Model of Climate Change, Energy Economics, Vol. 20, No. 5-6, December 1998.

Appendix: Models which explicitly account for uncertainty in the model structure:

Model	*Underlying Model Type*	*Type of Uncertainty Analysis*	*Which uncertainties*	*Key result*
AS/ExM	Policy evaluation, Simulation; Optimization	Exploratory modeling (explicitly comparing the performance of alternative strategies against a large number of plausible futures); Compares an adaptive strategy with two best-estimate policies.	• climate sensitivity • damages from temp. change • abatement cost reduction from technological innovation	Best-estimate policies incur large costs relative to optimum policy if state of the world turns out to be different from what was assumed by best-estimate policy. Choice between best-estimate policies is strongly dependent on society's expectations about all three uncertainties. Adaptive strategy, with ability to make midcourse corrections, performs better on average than best-estimate policies unless society is virtually certain that one best-estimate is correct.
Lempert, R. J., M. E. Schlesinger and S. C. Bankes, When we don't know the costs or the benefits: Adaptive Strategies for Abating Climate Change, Climatic Change, 1996.				
CETA	Optimization	Sequential decision making under uncertainty	• warming per CO2 doubling • level parameter in damage function • power parameter in damage function	If an optimal control policy is used, the benefit of resolving uncertainty is high, but resolving uncertainty now vs. in 20 years is not worth much. If an arbitrary political policy is used, and if resolving uncertainty now would imply that an optimal policy would be used, then there is a high premium on resolving uncertainty now vs. later.
Peck, Stephen C. and Thomas J. Teisberg, Global Warming Uncertainties and the Value of Information: An Analysis using CETA, Resource and Energy Economics 15, 1993.				
CETA-R	Optimization	Stochastic Losses (represented in the utility function)	• large loss	Results are very sensitive to the risk perception of the experts considered in the specification of the loss probability function. No significant control occurs before 2060, even with the high risk.
Peck, Stephen C and Thomas J. Teisberg, Optimal CO2 Control Policy with Stochastic Losses from Temperature Rise, Climatic Change, Vol. 31, 1995.				
HCRA	Policy evaluation	Sequential decision making under uncertainty	• climate sensitivity • economic costs of emissions reductions • climate target delta T.	Value uncertainties may be more salient for policy choice. Policies that are sequentially revised as new information becomes available may be superior.
Hammit, James K., Outcome and Value Uncertainties in Global Change Policy, 12/94.				
DIAM	Optimization	Sequential decision making under uncertainty	• stabilization limit (stochastic constraint) • impact costs	• Possibility of low levels of stabilization limits has large influence on optimal path. Even though this occurs with low probability, the large cost assigned to the constraint drives the outcome. • Consideration of impact costs leads to different time profiles than optimization under a stabilization constraint (fixed or stochastic.)
Grubb, Michael, Technologies, energy systems and the timing of CO2 emissions abatement, Energy Policy, Vol. 25, No.2, 1997.				
DICE	Optimization	Monte Carlo Analysis (using representative scenarios); Sequential decision making under uncertainty	• rate of population growth • productivity growth • discount rate • GHG-output ratio • damage function intercept • climate-GHG sensitivity • mitigation cost fctn. intercept • atmospheric detention rate	Carbon tax might be a more efficient instrument in light of enormous uncertainties. Carbon tax is more invariant across resolution of uncertainties than optimal GHG control rate. Value of Early Information can help understand how investments of resources to obtain better information about the future climate and social sciences pay off.
Nordhaus, W., Managing the Global Commons, MIT Press, Cambridge, 1994.				

FUND 1.5	Optimization/ Simulation	Monte Carlo analysis; Propagation of selected parameters.	selected parameters, including: • socio-economic drivers • carbon cycle/climate • climate change impacts • emissions reduction	• The business as usual scenario leads to an unbounded loss when uncertainty is included (though the divergence is slow). This does not occur with the emissions reduction scenarios. • Optimal emissions reduction is more strict under uncertainty than under certainty. • Under uncertainty, there is no emission trajectory that avoids risk of both severe costs of emission reductions and severe impacts of climate change.
Tol, Richard S. J., Tvd Burg, HMA Jansen, H. Verbruggen, The Climate Fund, Some notions on the socioeconomic impact of greenhouse gas emissions and emission reduction in an international context, Institute for Environmental Studies, RG5/03, Vrije Universiteit, Amsterdam, 1995. Tol, Richard S. J., A Decision-Analytic Treatise of the Enhanced Greenhouse Effect, Institute for Environmental Studies, RG5/03, Vrije Universiteit, Amsterdam, 1997.				
ICAM-2	Simulation; Various decision rules	Propagation of uncertainty	• parameters (up to 25) • decision rules and metrics • model structure	Optimal decision depends on the decision rule. None of the policies are stochastically dominant.
Dowlatabadi, Hadi and Matt Ball, An overview of the Integrated Climate Assessment Model Version 2, (ICAM-2), presented at the Western Economics Association Conference, 6/29/94.				
Kelly/ Kolstad	Stochastic infinite horizon optimization	Sequential decision making under uncertainty with endogenous and continuous learning about uncertain parameter.	• uncertainty about climate sensitivity coupled with random unobserved shock to temperature	Interplay between learning about the climate change problem and decisions to control the problem: It can take a very long time to resolve the uncertainty, during which significant suboptimal control can take place (relative to perfect information).
Kelly, David L., and Charles D. Kolstad, Tracking the Climate Change Footprint: Stochastic Learning About Climate Change, University of California Santa Barbara Economics Working Paper 3-96R (Nov. 1996).				
MERGE 2.0	Optimization	Sequential decision making under uncertainty	• high-damage and low-damage scenario	With a small chance of high damages, a hedging strategy departs only slightly from the low-damage case. Hedging strategy is sensitive to date at which uncertainty is resolved.
Manne, Alan and Richard Richels, The Greenhouse debate—Economic Efficiency, Burden Sharing and Hedging Strategies, April 1995.				
PAGE 95	Policy evaluation; Stochastic simulation.	Propagate uncertainty about input parameters through model; Partial Rank Coefficients between inputs and output.	80 uncertain parameters: • scientific • costs of control • costs of adaptation • valuation of impacts	Important factors come from all four groups of inputs to the model. Most important parameters are preventive costs of CO2 and temperature sensitivity.
Plambeck, Erica L. and Chris Hope, Page 95, An updated valuation of the impacts of global warming, Energy Policy, Vol. 24, No. 9, 1996.				
PEF	Policy evaluation; Deterministic models and decision tree.	Decision analysis	• (scenarios only ?)	As climate change becomes worse, more action is warranted Uncertainty in impacts may be as important as uncertainty in extent of climate change. Adaptation policies could have larger effects on impacts more quickly. Mitigation is less effective when adaptation is high, but not vice versa.
Cohan, David, Stafford, R., Scheraga, J., Herrod, S., The Global Climate Policy Evaluation Framework, A&WMA, 4/1994.				
PRICE	Optimization	5 different approaches to estimate Value of Information about uncertain parameters; Value of Early Information	8 uncertain parameters (same as in DICE); 5 states of the world.	Damages from Climate change and costs of reducing GHG emissions are most important. Resolving their uncertainties would contribute 75 % of the value of improved knowledge.

SLICE	Finite horizon stochastic program/ Optimization	Sequential decision making with continuous learning (rate of learning is exogenous parameter)	• climate damage	The irreversibility of investment capital has a stronger effect than irreversibilities in climate change. Thus uncertainty and learning tend to bias emission control downward relative to the case of uncertainty but no learning.
Kolstad, Charles D., Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions, JEEM, 31:1-18 (1996).				
TARGETS	Stochastic simulation	Cultural Theory; The parameters selected are those that serve as accessories for model routes used in cultural theory, while cultural perspectives serve to determine the distribution interval.	• CO_2 fertilization • soil moisture changes • migration of ecosystems • temp. feedback on vegetation • temp. feedback on production • sulphate aerosols • water vapor • clouds • policy measures	Takes into account a variety of perspectives in relation to uncertainty, which allows for rendering subjective judgment explicit.
van Asselt, Marjolein B.A. and Jan Rotmans, Uncertainty in Perspective, Global Environmental Change, Vol. 6 No. 2, 1996.				
YOHE/Connecticut	Optimization	Sequential decision making under uncertainty	• population growth • technological change in energy supply • depletion factor in fossil fuel price • interfuel elasticity of substitution • others that play less significant roles in the distribution of emissions	• Little or no emissions reduction is warranted over the near term even as a hedge against the possibility of having to meet severely binding concentration limits in the not too distant future. • Modest emissions reduction can be supported when hedging against high consequence/low probability events across a wide range of emissions futures • Hedging to achieve “tolerable windows” proposed by German Advisory Board on Climate Change would require significant near term emissions reduction at enormous cost.
Yohe, Gary and Rodney Wallace, Near term mitigation policy for global change under uncertainty: Minimizing the expected cost of meeting unknown concentration thresholds, Environmental Modeling and Assessment, Vol. 1, No 1,2, June 1996.				
Zapert et al. adaptation of IMAGE	Policy Evaluation	Propagation of uncertainty	• Initial state and/or stochastic noise are modeled for 155 uncertain parameters (mostly physical climate descriptors)	• Even conservative uncertainty estimates result in scenario overlap of several decades during which the consequences of any actions affecting the environment could be difficult to identify with sufficient level of confidence. • In general, the stochastic fluctuations contribute more to the uncertainty than the initial state measurements.
Zapert, R., Gaertner, P. S., Filar, J. A., Uncertainty Propagation within an Integrated Model of Climate Change, Energy Economics, Vol. 20, No. 5-6, December 1998.				