J. ENVIRONMENTAL SYSTEMS, Vol. 31(1) 1-21, 2004-2005

ESTIMATING PROBABILITY DISTRIBUTIONS FROM COMPLEX MODELS WITH BIFURCATIONS: THE CASE OF OCEAN CIRCULATION COLLAPSE*

MORT WEBSTER
Massachusetts Institute of Technology

JEFF SCOTT

ANDREI SOKOLOV

PETER STONE

Massachusetts Institute of Technology

ABSTRACT

Studying the uncertainty in computationally expensive models has required the development of specialized methods, including alternative sampling techniques and response surface approaches. However, existing techniques for response surface development break down when the model being studied exhibits discontinuities or bifurcations. One uncertain variable that exhibits this behavior is the thermohaline circulation (THC) as modeled in three dimensional general circulation models. This is a critical uncertainty for climate change policy studies. We investigate the development of a response surface for studying uncertainty in THC using the Deterministic Equivalent Modeling Method, a stochastic technique using expansions in orthogonal polynomials. We show that this approach is unable to reasonably approximate the model response. We demonstrate an alternative representation that accurately simulates the model's response, using a basis function with properties similar to the model's response over the uncertain parameter space. This indicates useful directions for future methodological improvements.

*This research was supported in part by the Methods and Models for Integrated Assessments Program of the National Science Foundation, Grant ATM-9909139; by the Office of Science (BER), U.S. Department of Energy, Grant Nos. DE-FG02-02ER63468 and DE-FG02-93ER61677; and by the MIT Joint Program on the Science and Policy of Global Change (JPSPGC). Financial support does not constitute an endorsement by NSF, DOE, or JPSPGC of the views expressed in this article.

INTRODUCTION

Estimating probability distributions of uncertain model outputs has long been a challenge for models requiring large amounts of computation time. A variety of methods have been developed for this problem, including specialized sampling methods [1] and constructing response surface approximation methods [2, 3]. One obstacle to using most response surface methods occurs when the model response exhibits discontinuities or bifurcations.

An example of bifurcating behavior is the change in the circulation of the North Atlantic Ocean in long-term climate change projections. The thermohaline circulation (THC), or more formally, the zonally averaged meridional overturning circulation (MOC), refers to the circulation pattern of the North Atlantic Ocean in which warm surface water from the tropics travels northward, considerably warming mid and high latitudes in the Northem Hemisphere around the globe. This circulation is driven by deep water formation in the northern North Atlantic near Greenland, which is caused by the water becoming colder until it reaches a critical density that causes it to sink. As a possible consequence of climate change, it is hypothesized that warmer temperatures and increased freshwater runoff could prevent the water from reaching its critical threshold density, thus shutting off this circulation.

The possibility of a collapse of the North Atlantic thermohaline circulation is one of the more severe potential impacts of climate change, and therefore is relevant to policy discussions [4]. A critical question, therefore, is: What is the probability of a THC collapse in the future?

One approach is to use simplified ocean models, which can reasonably be run for a large number of parametric assumptions [5]. However, for a more realistic representation of the ocean dynamics, one would ideally use a high-resolution three-dimensional ocean general circulation (GCM) model, coupled with a 3-D atmospheric GCM. A single simulation of several centuries with such models generally requires weeks to months on a supercomputer. Thus, even the small number of simulations (typically ~50 or more) required by methods such as Latin Hypercube Sampling [1] is prohibitive. Moreover, to inform policy, we need to know how the probability of a THC collapse will change with different policies, in addition to the reference case with no climate policy, requiring multiple sets of Monte Carlo simulations.

To obtain the desired information from the more detailed models, some kind of reduced-form response surface model is needed that replicates the full 3-D dynamic behavior of the ocean, yet is simple enough to perform Monte Carlo on to obtain probability estimates. However, commonly used methods do not apply to a system with a bifurcation, and ocean circulation models are well-known to exhibit exactly this kind of behavior.

In this article, we apply a commonly-used method for constructing optimal response-surface approximations for estimating the THC circulation from a 3-D ocean GCM. We will illustrate the challenges faced by this type of method, and demonstrate an alternative approach that is successful. The subsequent discussion frame directions for future research on more generalized approaches that can be applied to situations such as this one.

COUPLED CLIMATE MODEL DESCRIPTION

Our coupled model of intermediate complexity consists of a three-dimensional ocean GCM [6] coupled to a zonally-averaged, statistical-dynamical atmospheric model [7], and a thermodynamic sea-ice model [8]. Further detail on the general coupled model can be found in [9, 10].

Our model's open passage through our idealized "Canadian Archipelago" plays an important role in the increased CO_2 simulations. Previous studies have speculated on the sensitivity of the ocean circulation and climate to freshwater discharge into the Arctic basin and subsequent flow into the Northern Atlantic [11-14]. Our model employs a flexible river-routing scheme for anomalous runoff (as calculated in the atmospheric sub-component). In the southern hemisphere, for simplicity (and lacking a river network in this idealized topography) this runoff is distributed evenly over all ocean points. In the northern hemisphere, however, all anomalous runoff is diverted to the Arctic Ocean at 72-76°N between 96° and 260° in longitude. This diversion of anomalous runoff was necessary in order to achieve a complete collapse of the THC across a sizeable portion of our parameter phase space. Given this and other model idealizations, our model cannot be expected to give realistic information about when a collapse will occur. Rather, our goal is to study qualitatively how the collapse depends on the parameters. Such a study has previously only been carried out with two-dimensional models of the ocean basins [5].

For our climate change scenarios, the level of CO_2 is increased in the atmospheric model at a constant compounded rate for 100 years and then held constant at this resulting level. Thus, the rate of increase is proportional to the final change of forcing in the atmosphere. For the climate sensitivity parameter, different sensitivities are obtained by varying the strength of cloud feedback [7]. Varying the feedback allows the 2D atmospheric model to mimic the results of AGCMs with different sensitivities when coupled to a mixed layer ocean model, with a fixed ocean heat transport. Values of climate sensitivity shown throughout the article represent an equilibrium sensitivity of the atmospheric model coupled to a mixed layer ocean model for a doubling of CO_2 concentration. However, defined in this way, the climate sensitivity does not precisely match the climate sensitivity of the coupled climate model because of the interaction between the atmosphere and the dynamic ocean.

We explore the uncertainty in the maximum overturning in the North Atlantic that is a consequence of uncertainty in two critical characteristics of climate system: the climate sensitivity and the rate of increase of CO_2 forcing. These uncertainties have previously been identified as primary determinants in ocean circulation changes [4, 5]. The assumed distribution for climate sensitivity, defined as the equilibrium warming resulting from a doubling of CO_2 concentrations, comes from [15], and is derived by updating expert priors with constraints from 20th-century observations. The probability density function (PDF) of the rate of CO_2 increase, driven primarily by anthropogenic emissions, is taken from [16], and is calculated from a Monte Carlo analysis of a macroeconomic model with uncertainty in economic growth rates and rates of energy efficiency improvement. Both PDFs are shown in Figure 1. The CO_2 forcing rate of increase is applied for the first 100 years of the simulation, and then CO_2 concentrations are held constant for the remaining 900 years of the of the simulation.

ALTERNATIVE METHODS FOR ESTIMATING PROBABILITIES

Overview of Methods

This section reviews the alternative methods for obtaining the uncertainty in an outcome from a deterministic computational model. Most simulation models are sufficiently complex that direct analytical solutions are not an option. The standard approach for uncertainty propagation is Monte Carlo simulation [17, 18], in which random samples are drawn from probability distributions of input parameters, the model is simulated for each random draw, and the frequency distribution of model outcomes provide the estimate of the probability distribution. The challenge to applying Monte Carlo comes when a model cannot be simulated thousands or tens of thousands of times.

As discussed above, one approach is to use variance reduction methods for sampling from parameter distributions, so that fewer samples are needed for the estimated probability distribution of the outcome to converge. One popular and effective approach is stratified sampling, as in the Latin Hypercube Sampling (LHS) method [1, 19]. If the goal of the analysis is to estimate the probability of an extreme event, an alternative is to use Importance Sampling [20], which focuses on the low-probability region of interest. As mentioned previously, 3D ocean circulation models are likely to be too expensive for LHS, especially when separate Monte Carlo simulations must be performed for several different policy options.

The other broad approach to estimating uncertainty from a computationally-intensive model is to construct a reduced-form model of the full model that produces a good approximation of the original model response with significantly less computation time. Reduced-form models can be further divided into two classes: theory-based or structural models and response surface approximations. Theory-based reduced-form models [21, 22] are simpler mathematical

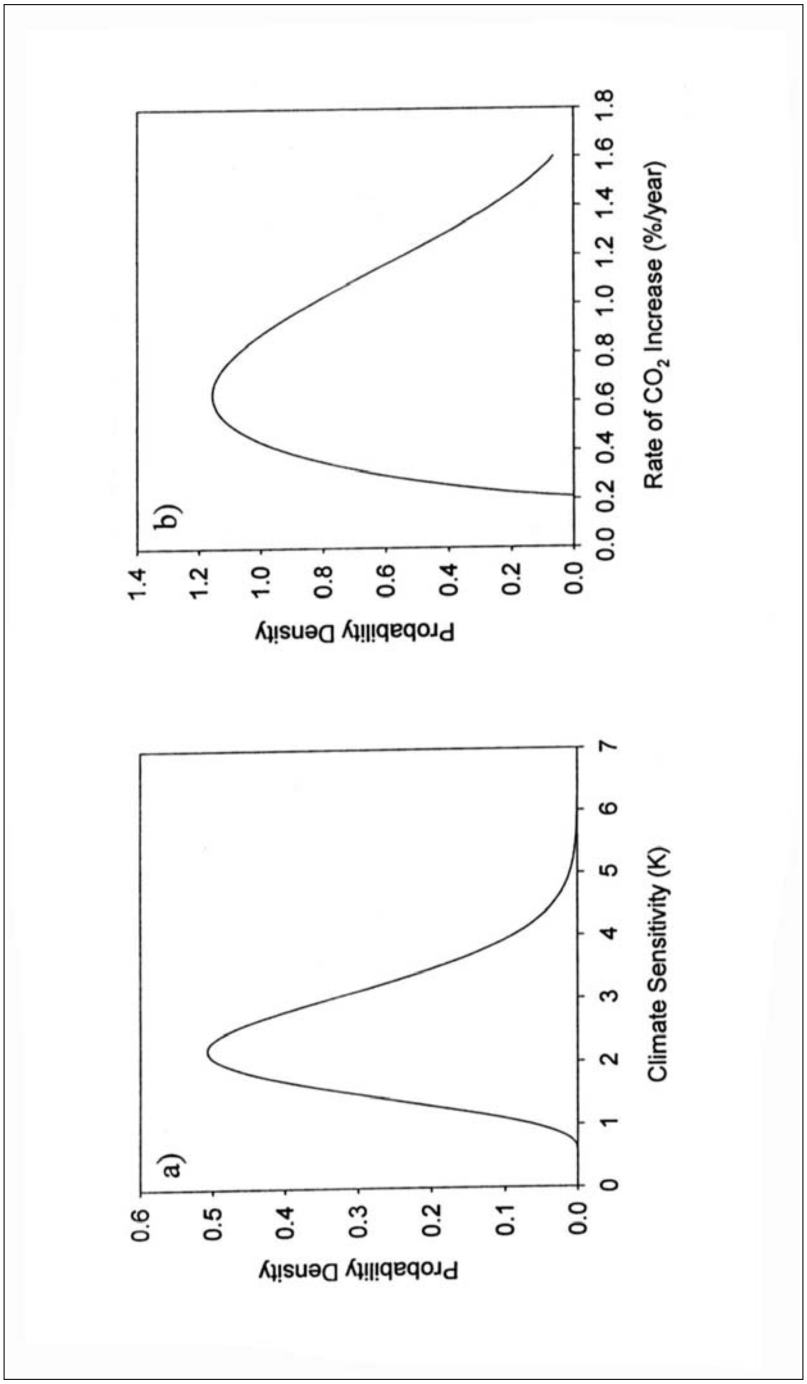

Figure 1. Probability distributions for uncertain parameters: a) climate sensitivity and b) rate of CO_2 forcing increase.

representations where the variables and equations still correspond to conceptual quantities and processes. This approach is primarily useful when transparency is critical for the reduced-form models behavior. The primary drawback is the extra time and effort required to develop a parsimonious closed-form model and the large number of runs of the original model to produce statistically acceptable parameter estimations.

he other subclass of reduced form models is response surface approximations. In these methods, a mathematical representation of the full model's response surface is developed, focusing only on the uncertain parameters for the particular analysis and their relationship to the model outcome(s) of interest. There is a variety of methods for response surface approximation, ranging from simple linear models to more sophisticated techniques. The choices that distinguish between these methods are:

1. The choice of the basis function, the fundamental elements in the equation(s) to be fitted to the model responses,
2. The choice of which parameter values to evaluate the full model at and use to fit, and
3. The choice of solution method, given a set of data points from the model and a set of coefficients to solve for in the fitted equation(s).

For example, standard linear approaches to response surface fitting [3] use first- or second-order polynomials of the uncertain parameters as a basis function, standard experimental design methods of choosing points for model evaluation, and minimize least-squared errors as the solution method to find the coefficients.

An alternative method for response surface approximation is the Deterministic Equivalent Modeling Method (DEMM) [23, 24]—this method is also sometimes referred to as the Probabilistic Collocation method (PCM). It is equivalent to the Stochastic Response Surface Method (SRSM) developed by [2]. DEMM seeks to characterize the probabilistic response of the uncertain model output as an expansion in orthogonal polynomials. We describe DEMM in more detail below.

There are several factors that determine which of the above methods is appropriate for any given situation, both the general class of approach (variance reduction vs. reduced-form model) and the particular choice (LHS vs. importance sampling). One important factor is the number of uncertain parameters under investigation. The number of simulations to obtain an accurate fit grows slowly for some methods (e.g., LHS) but expands rapidly for others (e.g., DEMM). Another critical factor is whether any prior information on the shape of the response within the range of uncertainty exists. Some methods are "black-box," no prior knowledge is required, while others (e.g., importance sampling) require some knowledge. DEMM is a good choice of method for estimating the uncertainty in the THC because: 1) it is black-box, requiring no prior knowledge of the shape; 2) the number of uncertain parameters is small (two); and 3) independent estimations of uncertainty are required for many different policy cases, which makes LHS infeasible.

The Deterministic Equivalent Modeling Method

Although any numerical computer model is itself deterministic, by positing uncertainty in a model parameter, the model's outputs become uncertain and thus can be thought of as a random variable. One useful representation for a random variable is an expansion of some family of orthogonal polynomials B_N (x) with weighting coefficients a_i:

$$y = a_0 B_0 + a_1 B_1(x) + a_2 B_2(x) + \ldots + a_N B_N(x)$$

where x is also a random variable of known distribution. Any family of orthogonal polynomials can be used, including Legendre, Laguerre, or Hermite. This expansion is sometimes referred to as a polynomial chaos expansion [25].

DEMM differs from the traditional approaches in all three steps that define a response-surface method. We first address the choice of the basis functions. Since a model output y is some function of its uncertain input parameter x, we can use information about the probability density of x to choose basis functions for the expansion. We can derive the set of orthogonal polynomials weighted by the density function of the parameter, according to the definition of orthogonal polynomials:

$$\int_x P(x) H_i(x) H_j(x) dx = C_i \delta_{ij} \tag{1}$$

where $\delta ij = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$

$H_i(x)$ and $H_j(x)$ are orthogonal polynomial functions of x of order I and j, $P(x)$ is some weighting function, and C_i is some constant (this constant is usually 1, and thus omitted, when the polynomials are normalized) in other words, the integral of the product of two orthogonal polynomials of different order is always 0. By using the probability density function of an input as the weighting function $P(x)$, a set of orthogonal polynomials can be derived recursively. (The zeroth order polynomial is always assumed to equal one.)

We next approach the method for estimating the weighting coefficients, a_i. There is a class of methods designed for solving this problem known as the methods of weighted residuals (MWR) [26]. The residual at any realization x_j of the random variable x, for some approximation $\hat{y}(x)$ of the function $y(x)$ is simply the difference:

$$R_N(\tilde{a}, x_j) = y(x_j) - \hat{y}(\tilde{a}, x_j)$$

where $R_N(\tilde{a}, x_j)$ is the residual for an N-term expansion with weighting coefficients $\tilde{a} = \{a_1, a_2, ..., a_N\}$.

In general, MWR solves for N coefficients by solving the N relations:

$$\int_0^1 R_N(\tilde{a}, x) W_j(x) dx = 0 \qquad j = 1,2\ldots, N \tag{2}$$

Alternative schemes for MWR differ by the choice of the form of the weighting function, $W_j(x)$. Commonly used schemes include the least squares method, which chooses $W_j(x)$ to be $\frac{\partial R_N}{\partial a_j}$, or Galerkin's method, which chooses $W_j(x)$ to be the derivatives of the approximation $\frac{\partial y_N}{\partial a_j}$. The difficulty with these schemes is that they require the explicit analytical form of the model in order to solve for the weighting coefficients. Because our goal is to approximate the uncertainty in a model output for any model, however complex, a method that allows the model to be treated as a "black-box" is preferable. This leads us to choose the collocation method, which uses the dirac delta function as the weighing function:

$$W_j(x) = \delta(x - x_j), \qquad j = 1,2,\ldots, N.$$

Since the integral of a function multiplied by a delta function is just the function evaluated at that point, solving (2) is equivalent to solving:

$$R_N(\tilde{a}, x_j) = 0, \qquad j = 1,2,\ldots, N \tag{3}$$

In other words, we simply solve for the set of a_j such that the approximation is exactly equal to the model at N points, and thus only require the model solution at N points and not the explicit model equations.

The final step in determining the polynomial chaos expansion to approximate the random variable is to choose the points x_j at which we evaluate the "true" model $y(x)$, in order to solve for the a_i using equation (3). For this step, we borrow from the technique of Gaussian Quadrature, which uses the summation of orthogonal polynomials multiplied by weighting coefficients to approximate the solution of an integral. In Gaussian Quadrature, the optimal choice of abscissas at which to evaluate the function being integrated are the N roots of the Nth order orthogonal polynomial B_N(x) [27]. Similarly in DEMM, to solve for the N coefficients in the expansion.

$$a_0 + a_1 B_1(x) + \ldots + a_{N-1} B_{N-1}(x),$$

we use the residual evaluated at the N roots of $B_N(x)$, the orthogonal polynomial one order higher than the highest order term.

For multiple uncertain parameters, N roots are generated for each parameter to use as possible sample values. However, not all possible permutations of the N values for each parameter will necessarily be needed, depending on the number of terms in the expansion. Rather than combine sample values randomly, as in Latin Hypercube, we can use the probability density functions of the parameters to order the N possible values by likelihood. Then sample sets are formed by choosing permutations in decreasing order of joint probability, until the required number of sets has been formed.

DEMM cannot find a sufficiently accurate approximation in every case. In particular, discontinuities in the response surface result in poor approximations. The approximation must be checked against model results at values of the uncertain inputs other than those used to solve for the coefficients. An optimal choice of points to check the approximation against the model is based on the roots of the next higher orthogonal polynomial than the one used to find points to solve at. The roots of the next higher order polynomial will always interleave the lower order roots [27], and so these will test the approximation at a maximal distance from the fit values while still spanning the highest probability regions. Moreover, if the expansion of order N results in an inaccurate fit, we already have the model results needed to solve the fit of order $N + 1$. Once the expansion for the probabilistic model response is solved and found to be reasonably accurate, the approximate probability density function of the response can be derived by applying Monte Carlo simulation to this expansion.

DEMM and similar methods have been used successfully to explore the uncertainty in a variety of scientific, engineering, and economic modeling applications [23, 24, 28-32]. For many models, DEMM estimates multiple characteristics of the response distribution more efficiently than either modified sampling or traditional response surface approximation methods. DEMM's approach of representing the PDF of the uncertain response as an expansion of underlying PDFs, and of using probabilistic information in choosing the sample points for fitting the expansion, enable more efficient approximation of the overall response distribution relative to other methods.

RESULTS

Behavior of Ocean Model as Climate Sensitivity and CO_2 Forcing Changes

The behavior of the maximum overturning for eight different parameter samples is shown in Figure 2. Note that for the first 100 years while CO_2 is increasing, the circulation slows in all cases, and does not collapse completely. But after several centuries the bifurcating behavior is apparent. For samples of either high climate sensitivity or rapid rate of CO_2 increase, ocean overturning continues to slow and shows no sign of rebounding within 1000 years. For samples with relatively low sensitivity and slow rate of CO_2 increase, the circulation recovers to close to present-day levels within a few centuries.

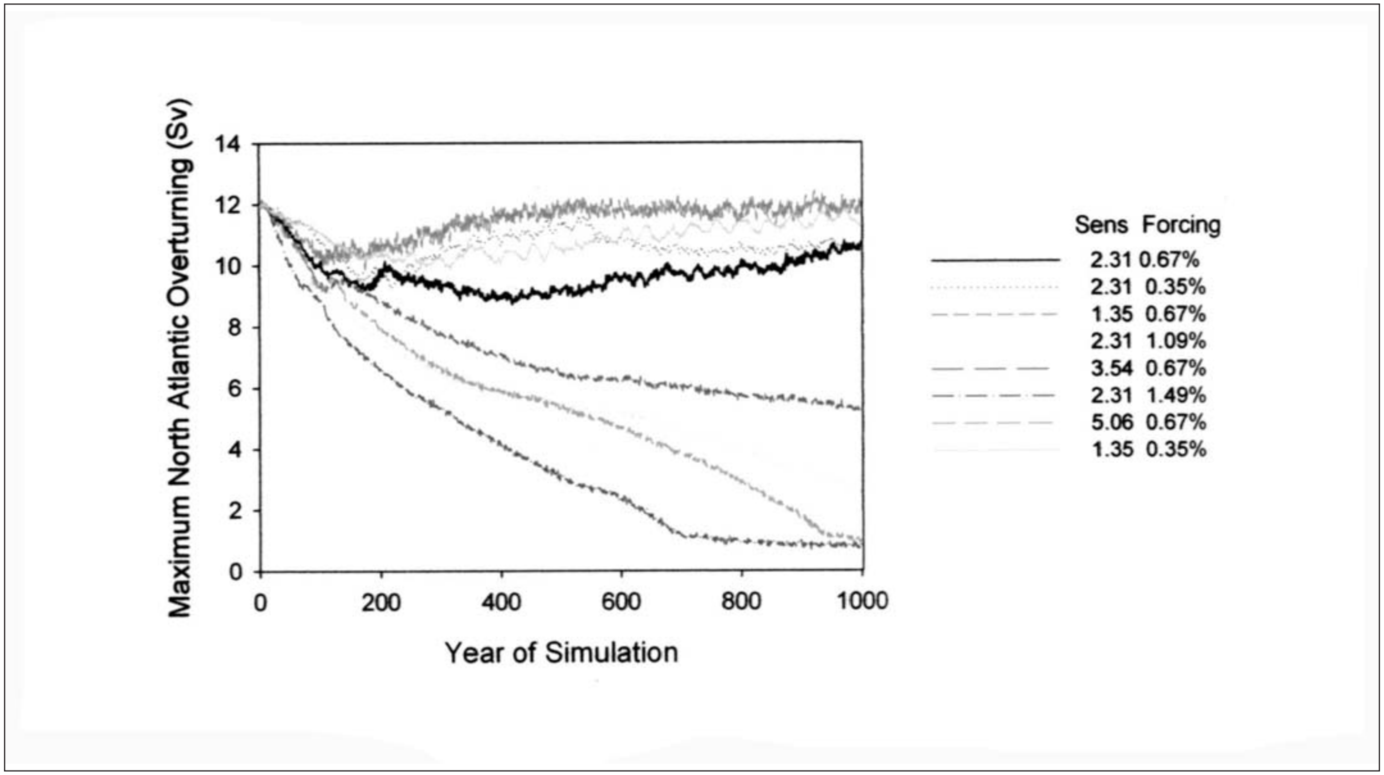

Figure 2. Time series of the maximum of the meridional overturning stream function in the North Atlantic for eight possible parameter sets.

Note that the transient behavior of the circulation in a simulation that does not recover (i.e., collapses) is continuous and smooth in the time dimension. The discontinuity is in the description of the circulation at one given point in time, for example in year 800, across all possible states of the world. The state of the circulation at some future time is the relevant outcome for policy studies.

Response Surface Fits with Different Methods

We first explore the application of DEMM to this problem. As described above, DEMM's use of orthogonal polynomials derived from the input PDFs is often superior to other response surface methods for non-linear surfaces, and has produced accurate estimates of probability distributions for a variety of applications including climate models.

Figure 3 shows the sample points in parameter space used to fit and test, respectively, a 3rd order DEMM approximation. This requires 8 simulations of the coupled model used to solve for the coefficients (circles) and additional 10

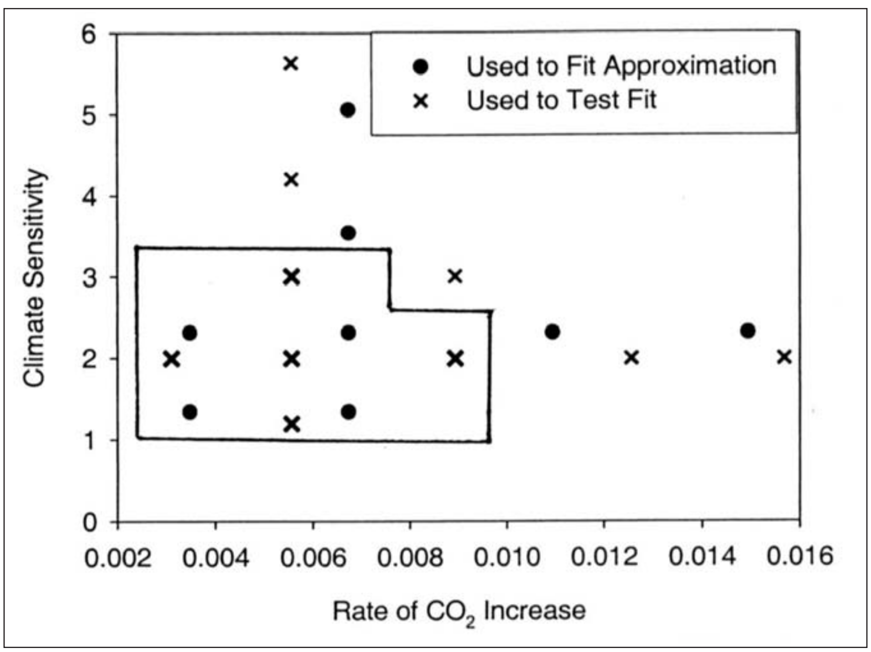

Figure 3. Initial parameter choices for fitting and testing DEMM approximation. The eight circles are parameter values used to fit the approximation, and the ten "x" symbols are used to compare the approximation to the actual model. Boxed symbols indicate parameter choices where the MOC recovers, and symbols not boxed indicate parameter choices where the MOC does not recover.

simulations used to test the goodness of fit (crosses). Note that the sample points are designed to optimally span the joint density function of the input parameters.

Before exploring response surfaces of ocean circulation strength, we first show the results for DEMM expansions of global mean surface air temperature (SAT) change. Third-order DEMM expansions for the parameter sets shown in Figure 3 result in approximations with sums of squared errors of less than 2% of the mean response value, accurately representing the response of the full climate model. Monte Carlo simulation is performed, drawing 10,000 random samples from the distributions for climate sensitivity and rate of CO_2 increase. The resulting PDFs of SAT change after 100 years and 1000 years are given in Figure 4.

Unfortunately, unlike surface air temperature change, the DEMM expansions for maximum North Atlantic overturning have unacceptably large errors for all years beyond year 200 (Figure 5). This is not surprising, as the surfaces span the discontinuity between the region where the overturning recovers and the region where it does not (see Figure 3).

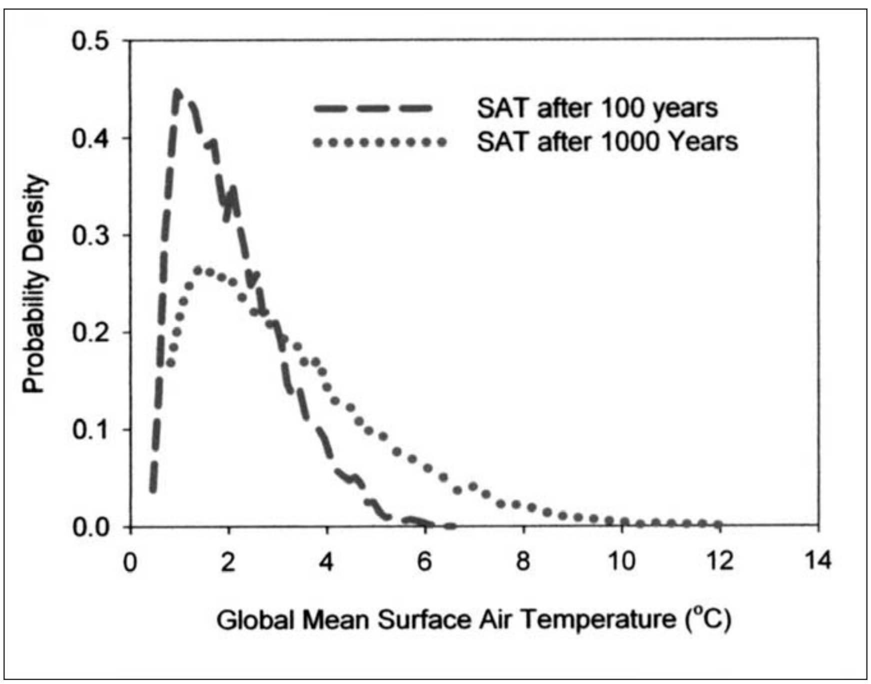

Figure 4. Estimated probability density functions for global mean surface air temperature after 100 years (dashed line) and 1000 years (dotted line).

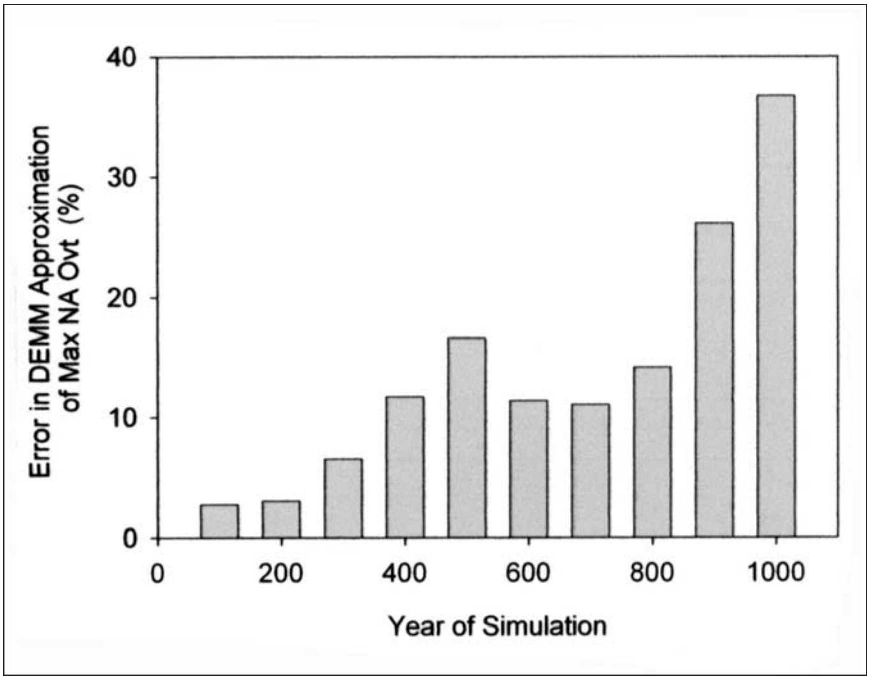

Figure 5. Errors in a 3rd-order DEMM expansion for the maximum overturning by century, measured as the average squared error relative to the mean value.

A second approach is to still use DEMM, but to fit it piecewise on either side of the discontinuity. This requires first that we identify the threshold between the region in parameter space where circulation recovers and the region where it does not. A total of 62 simulations were performed and used to calculate the critical threshold for circulation recovery. We find that the threshold is best identified by $s*r$, the product of the sensitivity and the rate of CO_2 increase (Figure 6). When $s*r < 1.72$, the circulation will recover, and when $s*r > 1.89$, the circulation collapses and does not recover within 1000 years.

Attempts to fit low-order DEMM approximations piecewise in each region of parameter space also fails to produce a reasonable representation of the ocean model's behavior. Figure 7 compares the best of the piecewise surfaces to the interpolated surface of the 62 GCM simulations. Monte Carlo simulations performed on DEMM approximations result in significant probability density for physically unrealistic values of maximum overturning below 0 and above 15 Sv. Further, piecewise fitting defeats the original purpose of selecting DEMM as a black-box method.

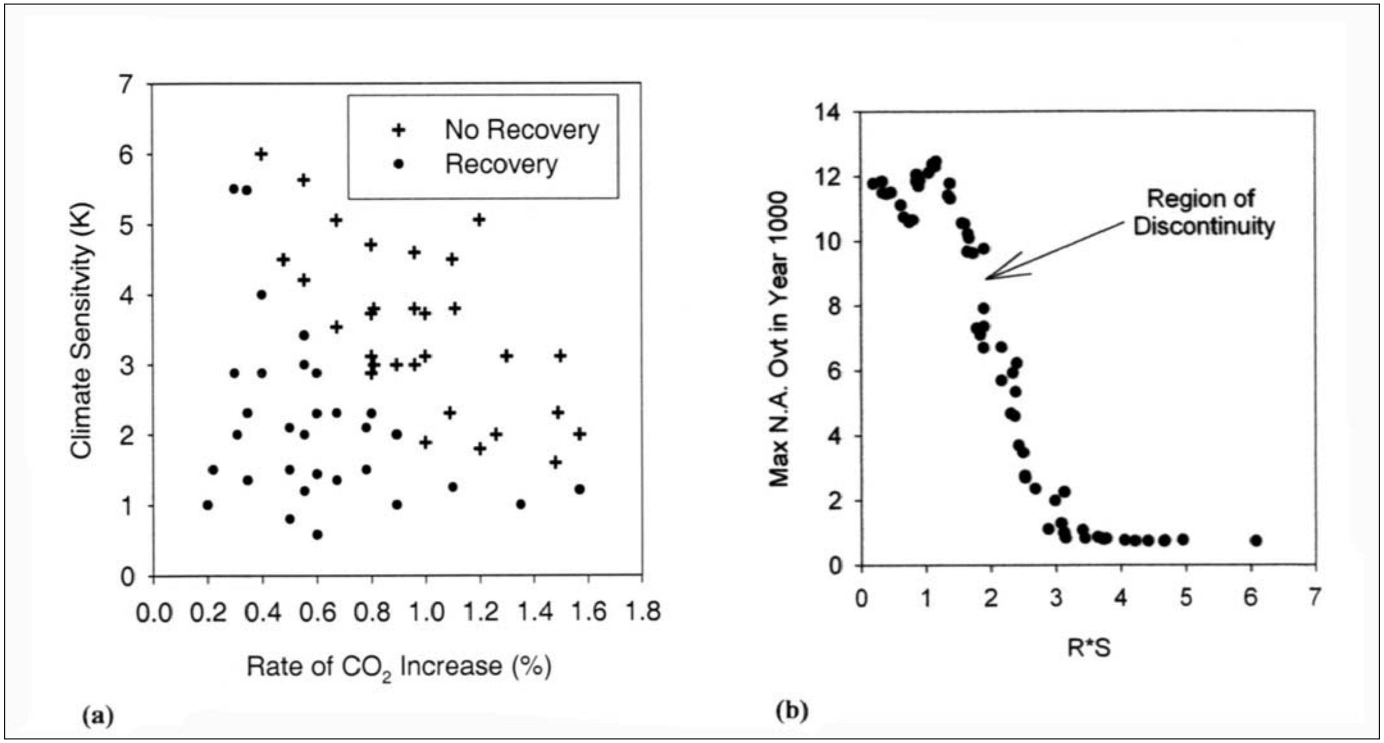

Figure 6. a) Parameter value pairs for all 62 simulations, circle points are values for which circulation recovers and plus points are values which collapse; b) Relationship between product of sensitivity and forcing rate with maximum overturning strength, gap indicates area of bifurcation.

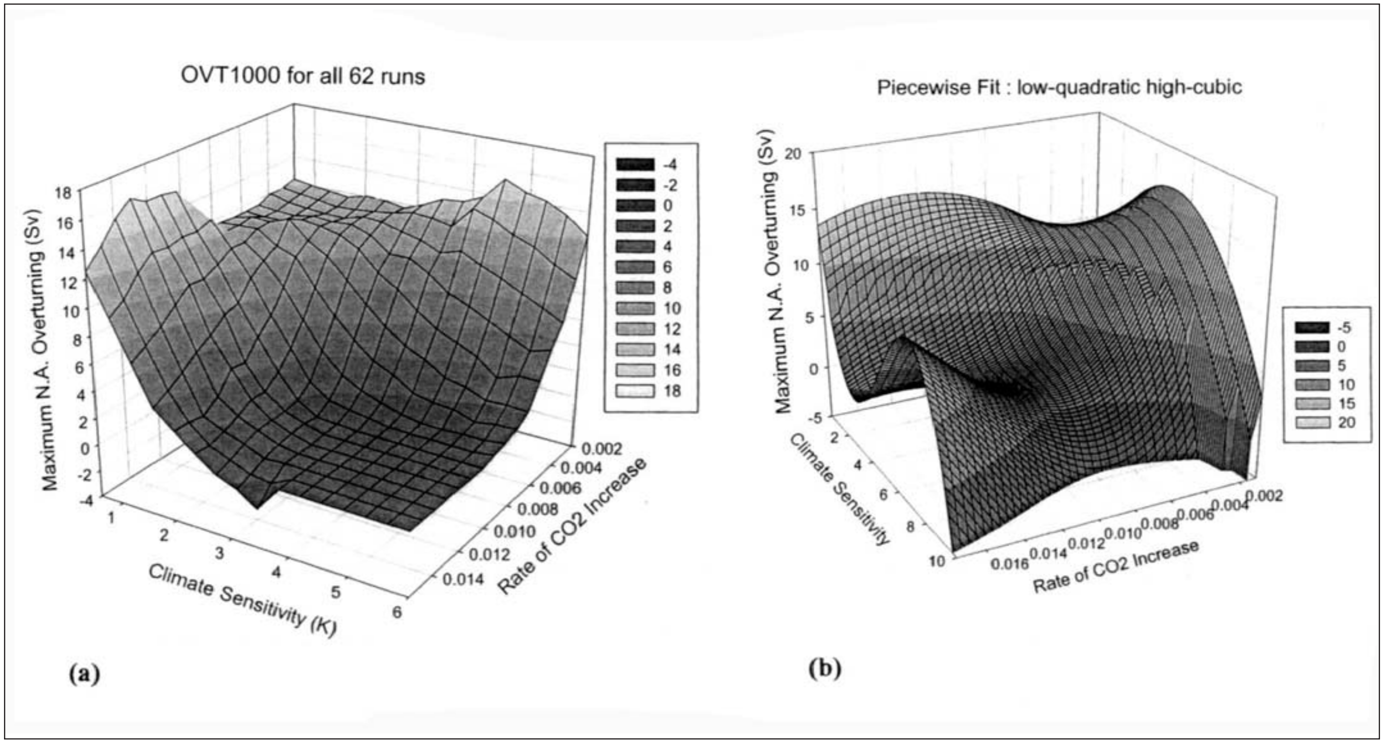

Figure 7. a) Interpolated response surface of maximum overturning in year 1000, using negative exponential smoothing over the 62 runs of the ocean model; b) a piecewise 3rd-order DEMM expansion fitting recovering and non-recovering regions separately.

A Successful Approximation Method

To understand why any polynomial-based approximation will fail to yield a reasonable fit to the model, consider the shape of the model's response surface in Figure 7a. Note that the overturning strength, when fully recovered, levels out at around 10-12 Sv. Similarly, overturning strength, once fully collapsed, levels out at close to 0 Sv. Thus, the projection into either sensitivity or CO_2 rate parameter space, the maximum overturning function has the shape of a logistical S-curve.

A low-order polynomial is unable to replicate this kind of S-curve shape, where function remains constant or approaches an asymptote above and below some critical values. As a demonstration, we apply DEMM to approximate the arctangent function, which exhibits this behavior. Treating arctan() as a black-box function, DEMM approximations are calculated, truncating terms at 3rd, 4th, 5th, and 7th order, respectively (Figure 8). Any low-order polynomial will have errors increasing exponentially in both directions beginning a short distance beyond the last model point used in the fit. A Monte Carlo with even low probability in these regions may yield large errors in the estimated PDF. Note that while a sufficiently

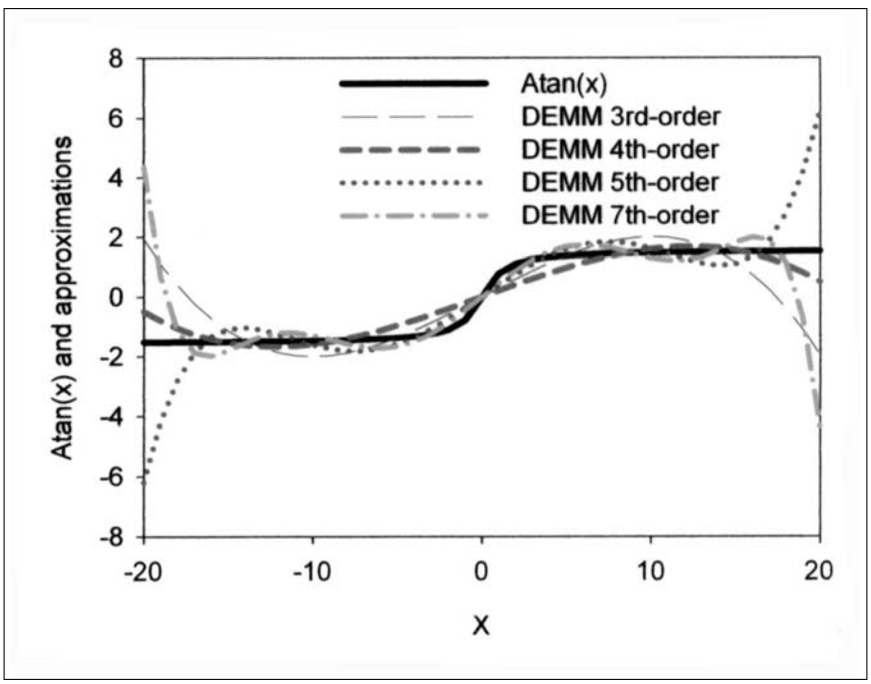

Figure 8. Arctan(x) (solid line) and DEMM expansions (dashed lines) of four different orders.

large number of expansion terms in orthogonal polynomials could be found that would reasonably approximate this kind of function, it would require even more model simulations than one would need to directly simulate with Latin Hypercube Sampling, and thus would yield no advantage.

The question becomes: is there an appropriate choice of basis function that WILL accurately replicate the model response across the parameter space? As described above, all response surface methods consist of a choice of basis function, a method of solving for coefficients, and a method of choosing points to evaluate the model for fitting. The problem here appears to be with the basis function choice. Having characterized the general shape of the response surface of the model, the ideal choice of basis function is one with the same logistical S-shaped curve. There are a number of functional forms with that shape from which to choose. One choice, from the example above, is the arctangent function.

We use the 62 simulations of the ocean model to fit the function

$$ovt = \beta - 2\pi\beta(\arctan[\alpha((s*r) - \delta)]) \quad (4)$$

where β is a shift parameter, α is an amplitude parameter, δ is the inflection point parameter, s is the climate sensitivity, and r is the rate of CO_2 increase in % per year. Thus, we need to solve for three free parameters, β, α, and δ, given a set of triplets (*s, r, ovt*). We solve for the parameters with ordinary least squares. The parameter values are given in Table 1.

Fitting this equation produces a response surface that very closely resembles Figure 7a, and has extremely small errors of at most a few percent (Table 1). We then perform Monte Carlo simulation on this approximation, drawing 10,000 random samples from the distributions of climate sensitivity and forcing rate. The resulting PDF of overturning for year 1000 is shown in Figure 9. To estimate the probability of a collapse, we note that all parameter choices that recover have maximum circulations of 9 Sv or greater, while parameter choices that do not recover have maximum circulations of 8 Sv or less (Figure 7b). By calculating the probability of a maximum overturning of 8 Sv or less, we estimate that the

Table 1. Parameters, Errors, and Estimated Probability of Circulation Collapse for Three Arctangent-Based Approximations of Maximum N.A. Circulation in Year 1000

# points used to fit	β	α	δ	Avg. squared error (Sv)	Avg. Abs. error (Sv)	Prob. of THC collapse
8	5.85	5.32	2.35	1.42	0.83	6.3%
18	6.13	2.59	2.21	0.69	0.62	11.6%
62	6.47	2.30	2.12	0.60	0.60	13.9%

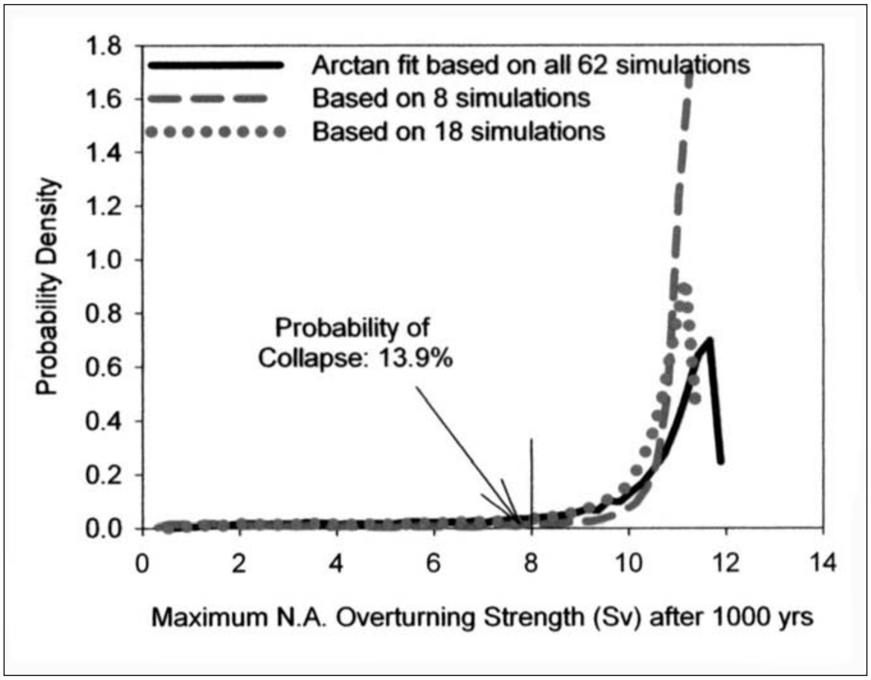

Figure 9. Probability distribution of the maximum North Atlantic overturning after 1000 years, based on approximation with arctangent basis function.

probability of a thermohaline circulation that collapses and does not recover within 1000 years is 13.9%.

This estimate is conditional on the assumed parameter distributions, but also importantly on the structural assumptions in the model. The true probability could be either higher or lower than this. More detailed studies are required with other coupled ocean-atmosphere GCMs for a range of assumptions to give better information on this likelihood.

While the fit with 62 simulations achieves an acceptable level of accuracy, the goal is to develop a method with far fewer simulations if possible. We develop two more fits of equation 4 using the points chosen for a 3rd-order DEMM expansion. The first uses only the 8 parameter sets used by DEMM to fit, and the second uses all 18 parameter sets from Figure 3 which consist of the points used by DEMM to fit and the points used to check the fit. The 8 point fit has larger errors, but the 18 point fit is nearly as accurate as the 62 point fit. The estimate of the probability of THC collapse from the 18 point fit is 11.6%, very close to the estimate from the 62 point fit. The results of a Monte Carlo on all three versions are shown in Figure 9.

DISCUSSION

In this study, we have attempted to find a way of approximating the response of a coupled ocean-atmosphere general circulation model to changes in two critical uncertainties: climate sensitivity and the rate of CO_2 increase. In particular, our interest is in describing the relationship between these parameters and the likelihood of a collapse of the thermohaline circulation in the North Atlantic. Because this response is discontinuous with a bifurcation, it poses a particular challenge to developing an accurate reduced-form that is amenable to multiple rounds of Monte Carlo simulation.

The solution to the methodological problem, while admittedly ad-hoc, points the way to new generalized techniques of response surface approximation. In the end, the obstacle to using existing methods was not so much the bifurcation, but the appropriate shape of the underlying basis functions. Although we leave the development of formal generalized methods to future work, needed improvements will be in the area of developing efficient methods for l) identifying the response surface shape characteristics; 2) choosing the appropriate basis functions for that shape, where the basis functions are chosen from a menu of options that include non-polynomial functions; and 3) identifying optimal points to sample the true model, given the choice of basis functions.

This study also suggests a useful general approach for policy-focused studies of uncertainty in climate change. There is a hierarchy of complexity for climate models, ranging from simple box and 1-D models, to earth models of intermediate complexity (EMICs) which are often 2-D or 3-D with limited resolution, to full 3-D GCMs. One way to use this spectrum of available tools in studying the uncertainty in any climate change process is to study the process with an EMIC, develop an appropriate basis function for a response surface, and then conduct limited simulations with a full GCM to fit the response surface. This would be a hybrid approach between a theory-based and a response surface reduced-form model.

REFERENCES

1. R. L. Iman and J. C. Helton, An Investigation of Uncertainty and Sensitivity Analysis Techniques for Computer Models, *Risk Analysis, 8*:1, pp. 71-90, 1988.
2. S. S. Isukapalli, A. Roy, and P. G. Georgopoulos, Stochastic Response Surface Methods (SRSMs) for Uncertainty Propagation: Application to Environmental and Biological Systems, *Risk Analysis, 18*:3, pp. 351-363, 1998.
3. G. E. P. Box and N. R. Draper, *Empirical Model-Building and Response Surfaces,* Wiley, New York, 1987.
4. K. Keller, B. M. Bolker, and D. F. Bradford, Uncertain Climate Thresholds and Optimal Economic Growth, *Journal of Environment Economy Management, 48,* pp. 723-741, 2004.

5. A. Schmittner and T. F. Stocker, The Stability of the Thermohaline Circulation in Global Warming Experiments, *Journal of Climatology, 12,* pp. 1117-1133, 1999.
6. J. Marshall, A. Adcroft, C. Hill, L. Perelman, and C. Heisey, A Finite-Volume, Incompressible Navier-Stokes Model for Studies of the Ocean on Parallel Computers, *Journal of Geophysical Resource, 102,* pp. 5753-5766,1997.
7. A. Sokolov and P. H. Stone, A Flexible Climate Model for Use in Integrated Assessments, *Climate Dynamics, 14,* pp. 291-303, 1998.
8. M. Winton, A Reformulated Three-Layer Sea Ice Model, *Journal of Atmosphere Ocean Technology, 17,* pp. 525-531, 2000.
9. S. Dutkiewicz, A. Sokolov, J. R. Scott, and P. H. Stone, A Three-Dimensional Ocean Sea-Ice-Carbon Cycle Model and Its Coupling to a Two-Dimensional Atmospheric Model: Uses in Climate Change Studies, *Joint Program on the Science and Policy of Global Change Report, 122,* MIT, Cambridge, Massachusetts, 2005.
10. J. R. Scott, A. P. Sokolov, and P. H. Stone, *Relative Roles of Climate Sensitivity and Forcing in Defining the Ocean Circulation Response to Climate Change,* 2006 (manuscript in preparation).
11. H. Goosse, T. Fichefet, and J.-M. Campin, The Effects of the Water Flow through the Canadian Archipelago in a Global Ice-Ocean Model, *Geophysical Research Letter, 24,*p p. 1507-1510, 1997.
12. B. J. Peterson, R. M. Holmes, J. W. McClelland, C. J. Vorosmarty, R. B. Lammers, A. I. Shiklomanov, I. A. Shiklomanov, and S. Rahmstorf, Increasing River Discharge into the Arctic Ocean, *Science, 298,* pp. 2171-2173, 2002.
13. M. Khodri, G. Ramstein, D. Paillard, J. C. Duplessy, M. Kageyama, and A. Ganopolski, Modelling the Climate Evolution from the Last Interglacial to the Start of the Last Glaciation: The Role of the Arctic Ocean Freshwater Budget, *Geophysical Research Letter, 30,* pp. 1606, doi:10.1029/2003GL017108, 2003.
14. P. Wu, R. Wood, and P. Stott, Human Influences on Increasing Arctic River Discharges, *Geophysical Research Letter, 32,* L02703, doi:10.1029/2004GL021570, 2005.
15. C. E. Forest, P. H. Stone, A. P. Sokolov, M. R. Allen, and M. Webster, Quantifying Uncertainties in Climate System Properties with the Use of Recent Climate Observations, *Science, 295,* pp. 113-117, 2002.
16. M. D. Webster, M. Babiker, M. Mayer, J. M. Reilly, J. Harnisch, M. C. Sarofim, and C. Wang, Uncertainty in Emissions Projections for Climate Models, *Atmosphere and Environment, 36,* pp. 3659-3670, 2002
17. J. M. Hammersley and D. C. Handscomb, *Monte Carlo Methods,* Wiley, New York, 1964.
18. M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods,* J. Wiley & Sons, New York, 1986.
19. M. D. McKay, R. J. Beckman, and W. J. Conover, A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code, *Technometrics, 21*:2, pp. 239-245, 1979.
20. E. Clark, Importance Sampling in Monte Carlo Analysis, *Operations Research, 9,* pp. 603-620, 1961.
21. W. D. Nordhaus and J. Boyer, *Warming the World: Economic Models of Global Warming,* MIT Press, Cambridge, Massachusetts, 2000.

22. L. J. Valverde, H. D. Jacoby, and G. M. Kaufman, Sequential Climate Decisions under Uncertainty: An Integrated Framework, *Environmental Modeling and Assessment, 4,* pp. 87-101, 1999.
23. M. A. Tatang, W. Pan, R. G. Prinn, and G. J. McRae, An Efficient Method for Parametric Uncertainty Analysis of Numerical Geophysical Models, *Journal of Geophysical Research, 102*:D18, p. 21, 1997.
24. M. D. Webster and A. P. Sokolov, A Methodology for Quantifying Uncertainty in Climate Projections, *Climatic Change, 46*:4, pp. 417-446, 2000.
25. N. Weiner, The Homogenous Chaos, *American Journal of Mathematics, 60,* pp. 897-936, 1938.
26. J. Villadsen and M. L. Michelsen, *Solution of Differential Equation Models by Polynomial Approximation,* Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1978.
27. W. H. Press, S. A. Teukolsky, W. V. Vetterling, and B. P. Flannery, *Numerical Recipes in C,* Cambridge University Press, Cambridge and New York, 1992.
28. W. Pan, M. A. Tatang, and G. J. McRae, Uncertainty Analysis of Indirect Radiative Forcing by Anthropogenic Sulfate Aerosols, *Journal of Geophysical Research, 103*:D4, p. 3815, 1998.
29. J. Calbo, W. Pan, M. D. Webster, G. J. McRae, and R. Primm, Parameterization of Urban Subgrid Scale Processes in Global Atmospheric Chemistry Models, *Journal of Geophysical Research, 103*:D3, p. 3437, 1998.
30. S. Balakrishnan, A. Roy, G. Ierapetritou, G. P. Flach, and P. G. Georgopolous, A Comparative Assessment of Efficient Uncertainty Analysis Techniques for Environmental Fate and Transport Models: Application to the FACT Model, *Journal of Hydrology, 307,* pp. 204-218, 2005.
31. F. Hossain, E. N. Anagnostou, and K.-H. Lee, A Non-Linear and Stochastic Response Surface Method for Bayesian Estimation of Uncertainty in Soil Moisture Simulation from a Land-Surface Model, *Nonlinear Process in Geophysics, 11,* pp. 427-440, 2004.
32. S. S. Isakapalli, A. Roy, A., and G. P. Georgopoulos, Efficient Sensitivity/Uncertainty Analysis Using the Combined Stochastic Response Surface Method and Automated Differentiation: Application to Environmental and Biological Systems, *Risk Analysis, 20*:5, pp. 591-602, 2000.

Direct reprint requests to:

Mort Webster
Visiting Professor
Earth Atmosphere and Planetary
Sciences Department
MIT Joint Program on the Science
and Policy of Global Change
Massachusetts Institute of Technology
E40-408
77 Massachusetts Avenue
Cambridge, MA 02139
e-mail: mort@mit.edu