CHAPTER 7

Towards the Probability of Rapid Climate Change

Peter G. Challenor, Robin K.S. Hankin and Robert Marsh
National Oceanography Centre, Southampton, University of Southampton, Southampton, Hampshire, UK

ABSTRACT: The climate of North West Europe is mild compared to Alaska because the overturning circulation in the Atlantic carries heat northwards. If this circulation were to collapse, as it appears to have done in the past, the climate of Europe, and the whole Northern Hemisphere, could change rapidly. This event is normally classified as a 'low probability/high impact' event, but there have been few attempts to quantify the probability. We present a statistical method that can be used, with a climate model, to estimate the probability of such a rapid climate change. To illustrate the method we use an intermediate complexity climate model, C-GOLDSTEIN combined with the SRES illustrative emission scenarios. The resulting probabilities are much higher than would be expected for a low probability event, around 30–40% depending upon the scenario. The most probable reason for this is the simplicity of the climate model, but the possibility exists that we may be at greater risk than we believed.

7.1 Introduction

Northwest Europe is up to 10°C warmer than equivalent latitudes in North America because a vigorous thermohaline circulation transports warm water northwards in the Atlantic basin (Rind et al., 1986). However, due to increasing concentrations of CO_2 in the atmosphere, this circulation could slow markedly (Cubasch et al., 2001) or even collapse (Rahmstorf and Ganopolski, 1999). The climatic impact of such a change in the ocean circulation would be severe, especially in Europe (Vellinga and Wood, 2002), but with worldwide consequences, and could happen on a rapid time scale. It is important therefore that we assess the risk of such a collapse in the thermohaline circulation (Marotzke, 2000). Recent studies have developed and adopted a probabilistic approach to address the climate response to rising levels of greenhouse gases (Wigley and Raper, 2001; Allen and Stainforth, 2002; Stainforth et al., 2005). However, to our knowledge, no study has yet addressed the probability of substantial weakening of the overturning circulation and the implied rapid climate change. In this paper we present a statistical technique that can be used to estimate the probability of such a rapid climate change using a model of the climate and illustrate it with a model of intermediate complexity.

7.2 A Method for Calculating Probabilities of Climate Events

Most modern climate models are deterministic: given a set of inputs they always give the same results on a given hardware platform. There are two standard ways to introduce an element of randomness and hence to make probabilistic predictions. The first is to use the internal, chaotic variability of the model. The initial conditions are varied by a small amount and an ensemble of model runs is performed. This method is widely used in weather forecasting. This is suitable for problems where the initial conditions are the important factor for predictability, predictability of the first kind. However, for long-range climate forecasting we believe we have predictability of the second kind where it is the boundary conditions that matter. In this case the perturbations need to be made on the boundary conditions. In our case these are the model parameters. A numerical model of the climate system contains a number of parameters, the 'true' value of which is unknown. If we represent our ignorance of these parameters in probabilistic terms we can propagate this uncertainty through the numerical model and hence produce a probability density function of the model outputs. This is the method we will use in this paper.

In essence, our method is to sample from a specified uncertainty distribution for the model input parameters, run the model for this combination of inputs and compute the output. This process is repeated many thousands of times to build up a Monte Carlo estimate of the probability density of the output. This type of Monte Carlo method is too computationally expensive for practical use; even intermediate complexity climate models such as C-GOLDSTEIN (Edwards and Marsh, 2005) are not fast enough to allow us to carry out such calculations with the required degree of accuracy. To overcome this problem we introduce the concept of an *emulator*. An emulator is a technique in which Bayesian statistical analysis is used to furnish a statistical approximation to the full dynamical model. In preference to a neural network (Knutti et al., 2003), we follow Oakley and O'Hagan (2002) and use a Gaussian process to build our emulator.

This has the advantage that is easier to understand and interpret, and every prediction comes with an associated uncertainty estimate. This means that the technique can reveal where the underlying assumptions are good and where they are not. Our emulators run about five orders of magnitude faster than a model such as C-GOLDSTEIN.

Full mathematical details of Gaussian processes and the Bayesian methods we use to fit them to the data are given in Oakley and O'Hagan (2002). The basic process of constructing and using an emulator is as follows:

1. For each of the parameters of the model, specify an uncertainty distribution (a 'prior') by expert elicitation and thereby define a prior pdf for the parameter space of the model.
2. We generate a set of parameter values that allow us to span the parameter space of these prior pdfs and run the climate model at each of these points to provide a calibration dataset of predicted MOC strength.
3. Estimate the parameters of the emulator using the calibration dataset using the methods given in Oakley and O'Hagan (2002).
4. Sample a large number (thousands) of points from the prior pdf.
5. Evaluate the emulator at each of these points. The output from the emulator then gives us an estimate of pdf of the variable being emulated from which we can calculate statistics such as the probability of being less than a specified value.

Ideally, in step 2 we would use an ensemble of model runs that spanned the complete parameter space of the model. However, as dimensionality increases this becomes difficult, and a factorial design soon requires an impractically large number of model runs. We therefore use the latin hypercube design (McKay et al., 1979), which requires us to specify in advance the number of model runs we can afford, in our example below this is 100. The range of each parameter is split up into this number of intervals of equal probability according to the uncertainty distribution of the input parameters. Our experience is that this distribution should be longer tailed than the input distribution used for the Monte Carlo calculations: the emulator is, along with all such estimation techniques, poor at extrapolation but good at interpolation so we want model runs out in the tails of the distribution to minimise the amount of extrapolation the emulator is called upon to do. For step 4, the order of the values of each parameter is now shuffled so that there is one and only one value in each of the equiprobable interval of each parameter (that is, the marginal distribution is unchanged), but the points are randomly scattered across multi-dimensional parameter space.

A Gaussian process is the extension of a multivariate Gaussian distribution to infinite dimension. For full mathematical details of Gaussian processes and the Bayesian methods we use to fit them to the data see Oakley and O'Hagan (2002). A Gaussian process is given by the sum of two terms: a deterministic, or mean, part and a stochastic part. The mean part can be considered as a general trend while the stochastic part is a local adjustment to the data. There is a trade-off between the variation explained by the mean function and the stochastic part. Following Oakley and O'Hagan (op. cit) we specify *a priori* that the mean function has a simple form (linear, in our case) with unknown parameters. The stochastic term in the Gaussian process is specified in terms of a correlation function. We use a Gaussian shape for the correlation function. This is parameterised by a correlation matrix. The elements of this matrix give the smoothness of the resulting Gaussian process. For simplicity we use a diagonal matrix, setting the off-diagonal terms to zero. These correlation scales cannot be estimated in a fully Bayesian way so are estimated using cross-validation. An alternative approach is to use regression techniques to model the mean function in a complex way. This means that the stochastic term is much less important and may make problems such as non-stationarity less important; for a non-climate example where this is done see Craig et al. (2001). Gaussian process emulators specified in this way are perfect interpolators of the data and it can be shown that any smooth function can be expressed as a Gaussian process.

It is important to specify the uncertainty distributions of the model inputs/parameters in step 2 carefully. In our case we elicit the information from experts, in this case the model builders and tuners. Our method was to request reasonable lower and upper limits for each parameter and interpret these as fifth and ninety-fifth percentiles of a log normal distribution. Because of the importance of the input distributions a sensitivity analysis was carried out to identify important input parameters; step 4 was repeated with doubled standard deviation for those parameters (see below for details). It is difficult to elicit the full joint input distribution so we have elicited the marginals and assumed that the inputs are independent. This assumption is almost certainly wrong and needs to be tested in further work. More complex elicitation methods (see the review by Garthwaite et al., 2005) need to be considered.

7.3 An Illustration: Emulating the MOC Response to Future CO_2 Forcing in C-GOLDSTEIN

To illustrate the methods described above we estimate the probability of the collapse of the thermohaline circulation under various emission scenarios using an intermediate complexity climate model. The climate model we use is C-GOLDSTEIN (Edwards and Marsh, 2005). This is a global model comprising a 3-D frictional-geostrophic ocean component configured in realistic geometry, including bathymetry, coupled to an energy-moisture balance model of the atmosphere and a thermodynamic model of sea ice. We use *a priori* independent log-normal distributions for 17 model parameters (Table 7.1). For 12 of the parameters, we use the distributions derived in an objective

Table 7.1 Mean value and standard deviation for each model parameter.

Parameter*	Mean	St. Dev.
Windstress scaling factor	1.734	0.1080
Ocean horizontal diffusivity (m^2s^{-1})	4342	437.9
Ocean vertical diffusivity (m^2s^{-1})	5.811e−05	1.428e−06
Ocean drag coefficient ($10^{-5}s^{-1}$)	3.625	0.3841
Atmospheric heat diffusivity (m^2s^{-1})	3.898e+06	2.705e+05
Atmospheric moisture diffusivity (m^2s^{-1})	1.631e+06	7.904e+04
'Width' of atmospheric heat diffusivity profile (radians)	1.347	0.1086
Slope (south-to-north) of atmospheric heat diffusivity profile	0.2178	0.04215
Zonal heat advection factor	0.1594	0.02254
Zonal moisture advection factor	0.1594	0.02254
Sea ice diffusivity (m^2s^{-1})	6786.0	831.6
Scaling factor for Atlantic-Pacific moisture flux (x 0.32 Sv)	0.9208	0.05056
Threshold humidity, for precipitation (%)	0.8511	0.01342
'Climate sensitivity'† (CO_2 radiative forcing, Wm^{-2})	6.000	5.000
Solar constant (Wm^{-2})	1368	3.755
Carbon removal e-folding time (years)	111.4	15.10
Greenland melt rate due to global warming‡ (Sv/°C)	0.01(Low) 0.03617 (High)	0.005793

*The first 15 parameters control the background model state. The first 12 of these have been objectively tuned in a previous study, while the last three (threshold humidity, climate sensitivity and solar output) are specified according to expert elicitation. The last two parameters control transient forcing (CO_2 concentration and ice sheet melting). Italics show the parameters that exert particular control on the strength of the overturning and which we varied in our experiment. For these parameters, the standard deviation was doubled in the cases with high uncertainty.

† The climate sensitivity parameter, ΔF_{2x}, determines an additional component in the outgoing planetary long-wave radiation according to $\Delta F_{2x}\ln(C/350)$, where C is the atmospheric concentration of CO_2 (units ppm). Values for ΔF_{2x} of 1, 6 and 11 Wm^{-2} yield 'orthodox' climate sensitivities of global-mean temperature rise under doubled CO_2 of around 0.5, 3.0 and 5.5 K, respectively.

‡ We used two mean values of the Greenland melt rate parameter (see main text).

tuning exercise (Hargreaves et al., 2004). For the others we elicited values from one of the model authors (Marsh) using the method described above. We specify particularly high variance for climate sensitivity, in line with recent results (Stainforth et al., 2005). We thus account for uncertainty in the model parameters, but not in the model physics (so called 'structural' uncertainty).

To generate our emulator as described above we need an ensemble of model runs to act as our 'training set'. We use an ensemble of 100 members in a latin hypercube design. We first 'spin up' the climate model for 4000 years to the present day (the year 2000, henceforth 'present day') in an ensemble of 100 members that coarsely samples from a range of values for fifteen key model parameters (see Table 7.1); the remaining two parameters are only used for simulations beyond the present day. Following 3800 years of spin-up under pre-industrial CO_2 concentration, the overturning reaches a near-equilibrium state in all ensemble members (see Figure 7.1). For the last 200 years of the spin-up, we specify historical CO_2 concentrations (Johnston, 2004), leading to slight (up to 5%) weakening in the overturning circulation. After the complete 4000-year spin-up we have 100 simulations of the current climate and the thermohaline circulation. Figure 7.2 shows fields of mean and standard deviation in surface temperature. The mean temperature field is similar to the ensemble–mean obtained by Hargreaves et al. (2005). The standard deviations reveal highest sensitivity to model parameters at high latitudes, especially in the northern hemisphere, principally due to differences (between ensemble members) in Arctic sea ice extent. We obtain an ensemble of present day overturning states, with ψ_{max} in the plausible range 12–23 Sv for 91 of the ensemble members (see Figure 7.1). The overturning circulation collapsed in the remaining nine members after the first 1000 years. Since we know that the overturning is not currently collapsed, we remove these from further analysis. This is a controversial point that we will return to in the discussion. We then specify future anthropogenic CO_2 emissions according to each of the six illustrative SRES scenarios (Nakicenovic and Swart, 2000) (A1B, A2, B1, B2, A1FI, A1T), to extend those simulations with a plausible overturning to the year 2100.

In extending the simulations over 2000–2100, we specify the SRES CO_2 emissions scenarios and introduce two further parameters (the last two parameters in Table 7.1) that relate to future melting of the Greenland ice sheet and the rate at which natural processes remove anthropogenic CO_2 from the atmosphere. The rate of CO_2 uptake is parameterised according to an e-folding timescale that represents the background absorption of excess CO_2 into marine and terrestrial reservoirs. This timescale can be

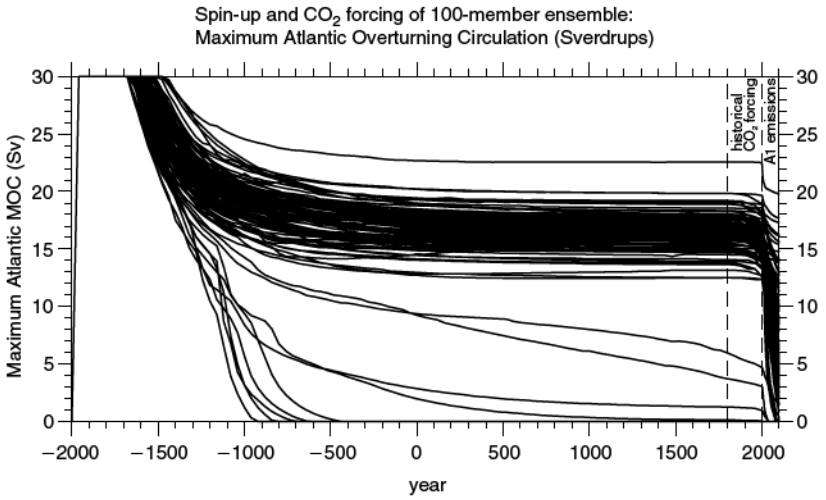

Figure 7.1 Spin-up of the Atlantic MOC, including CO_2 forcing from 1800.

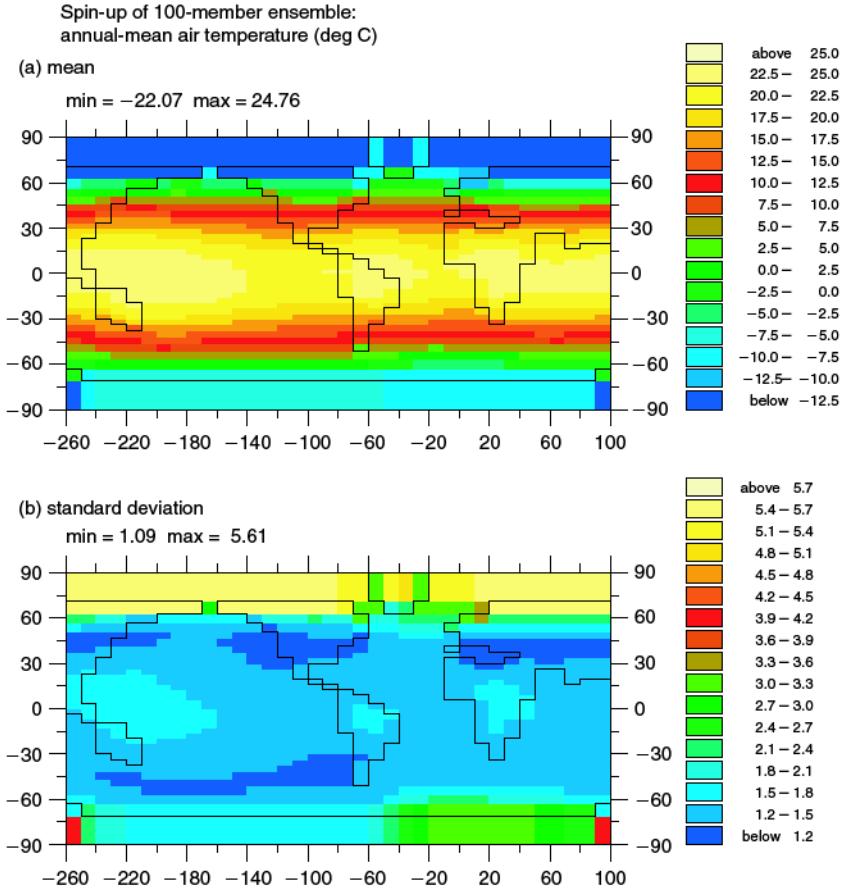

Figure 7.2 Mean and standard deviation of surface air temperature at year 2000.

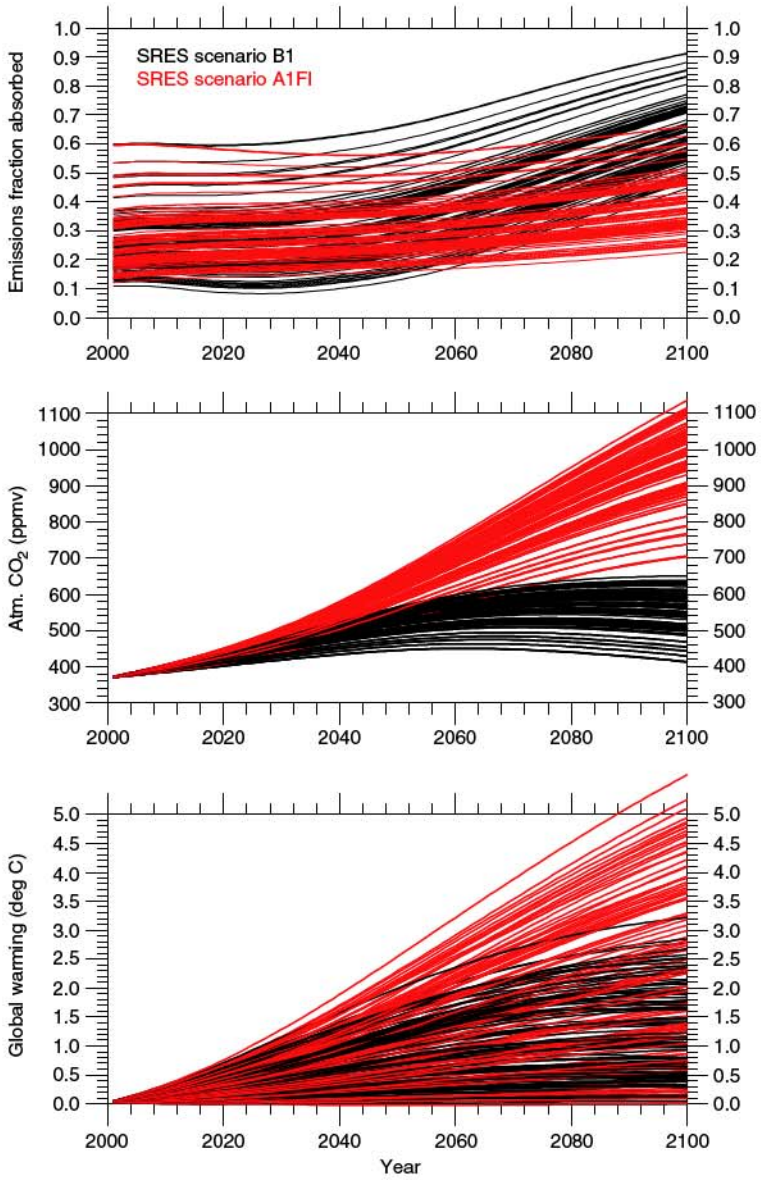

Figure 7.3 Time series of emitted CO_2 uptake, atmospheric CO_2 concentration and temperature rise over 2000–2100, under scenarios B1 and A1FI.

roughly equated with a fractional annual uptake of emissions. Timescales of 50, 100 and 300 years equate to fractional uptakes of around 50%, 30% and 10% respectively (see Figure 7.3, top panel), spanning the range of uncertainty in present and future uptake (Prentice et al., 2001). For each emissions scenario, a wide range of CO_2 rise is obtained, according to the uptake timescale (see Figure 7.3, middle panel). This in turn leads to a wide range of global-mean temperature rise, which is further broadened by the uncertainty in climate sensitivity (see Figure 7.3, bottom panel). The freshwater flux due to melting of the Greenland ice sheet is linearly proportional to the air temperature anomaly relative to 2000 (Rahmstorf and Ganopolski, 1999). This is consistent with evidence that the Greenland mass balance has only recently started changing (Bøggild et al. 2004). Over the range chosen for this parameter (combined with the uncertainty in emissions and climate sensitivity), the resultant melting equates to sea level rise by 2100 mostly in the range 0–30 cm (see Figure 7.4), consistent with predictions obtained with a complex ice sheet model (Huybrechts and de Wolde, 1999).

As a consequence of the applied forcing, ψ_{max} declines to varying degrees, in the range 10–90% in the case of the

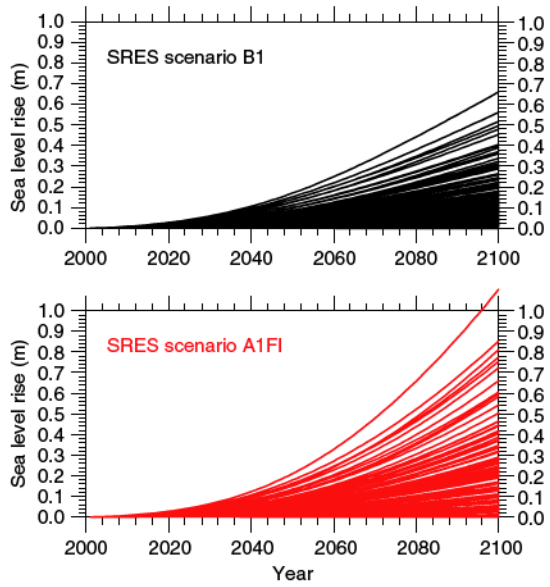

Figure 7.4 Sea level rise due to Greenland melting over 2000–2100, under scenarios B1 and A1FI.

A1FI scenario (see Figure 7.1). The range of MOC weakening is compatible with that suggested by IPCC (2001) AOGCM results. At 2100, the IPCC AOGCMs cover a range of +2 to −14 Sv with 9 model runs. The range for our 91 run ensemble is −1 to −17 with 90% between −2 and −15. Under the B1 scenario, the regional impact of this MOC slow-down is a local cooling in the Atlantic (see Figure 7.5, upper panel), also the location of highest standard deviation (Figure 7.5, lower panel), due to wide variation in the extent of slow-down. In several extreme cases (not clear from the ensemble-mean temperature change) of substantial slow-down, North Atlantic cooling under B1 exceeds 5°C. Under the A1FI scenario, global warming is amplified and the effect of MOC slow-down is to locally cancel warming (Figure 7.6, upper panel), and highest standard deviations are found in the Arctic (Figure 7.6, lower panel) due to disappearance of Atlantic sector Arctic sea ice cover in some ensemble members.

Using the model results for each SRES scenario at 2100, we build a statistical model (emulator) of ψ_{max} as a function of the model parameters. A separate emulator is built for each emissions scenario. We then use these six emulators, coupled with probability densities of parameter uncertainty, to calculate the probability that ψ_{max} falls below 5 Sv by 2100 using Monte Carlo methods. We use a sample size of 20,000 for all our Monte Carlo calculations. An initial, one-at-a-time, sensitivity analysis shows that the four most important parameters are: (1) sensitivity to global warming of the Greenland Ice Sheet melt rate, providing a fresh water influx to the mid-latitude North Atlantic that tends to suppress the overturning; (2) the rate at which anthropogenic CO_2 is removed from the atmosphere; (3) climate sensitivity (i.e., the global warming per CO_2 forcing); (4) a specified Atlantic-to-Pacific net moisture flux which increases Atlantic surface salinity and helps to support strong overturning. We perform a number of experiments calculating the probability of substantial slow-down of the overturning under variations in the values of these parameters and their uncertainties.

For each SRES scenario, we show in Table 7.2 the probability of substantial reduction in Atlantic overturning for five uncertainty cases. Each case is split into low and high *mean* Greenland melt rate, as this has been previously identified as a particularly crucial factor in the thermohaline circulation response to CO_2 forcing (Rahmstorf and Ganopolski, 1999). The probabilities in Table 7.2 are much higher than expected: substantial weakening of the overturning circulation is generally assumed to be a 'low probability, high impact' event, although 'low probability' tends not to be defined in numerical terms. Our results show that the probability is in the range 0.30–0.46 (depending on the SRES scenario adopted and the uncertainty case): this could not reasonably be described as 'low'. Even with the relatively benign B2 scenario we obtain probabilities of order 0.30, while with the fossil fuel intensive A1FI we obtain even higher probabilities, up to a maximum of 0.46.

Our probabilities are clearly less sensitive to the uncertainty case than to the SRES scenario. Increasing the mean Greenland melt rate from 'low' to 'high' increases only slightly the chance of shutdown in the circulation, probably because even the low melt rate already exceeds a threshold value (for substantial weakening of the overturning rate). The dependence of probability on parameter uncertainty is unclear, but any increase in uncertainty will broaden the distribution of the overturning strength and should theoretically lead to a higher proportion less than 5 Sv. While in some cases this is reflected in a slightly higher probability under higher parameter uncertainty (as expected), in other cases the probabilities are slightly lower. By comparing estimates from our sample of 20,000 between sub-samples of size 1,000 we estimate the standard error of our probability estimates to be about 0.01. If we had simple binomial sampling we would expect a standard error of about 0.05. We believe this difference in error comes from the correlation between estimates of the output. How much of this correlation comes from C-GOLDSTEIN and how much from the emulation process needs to be investigated. These error estimates imply that most of the random variation in our estimates is due to uncertainty coming from the fact that our emulation is not perfect, although some may also be caused by complex positive and negative feedbacks in the climate model.

7.4 Conclusions and Discussion

We have described a method that can be used to estimate the probability of a substantial slow-down in the Atlantic thermohaline circulation and a consequent rapid climate

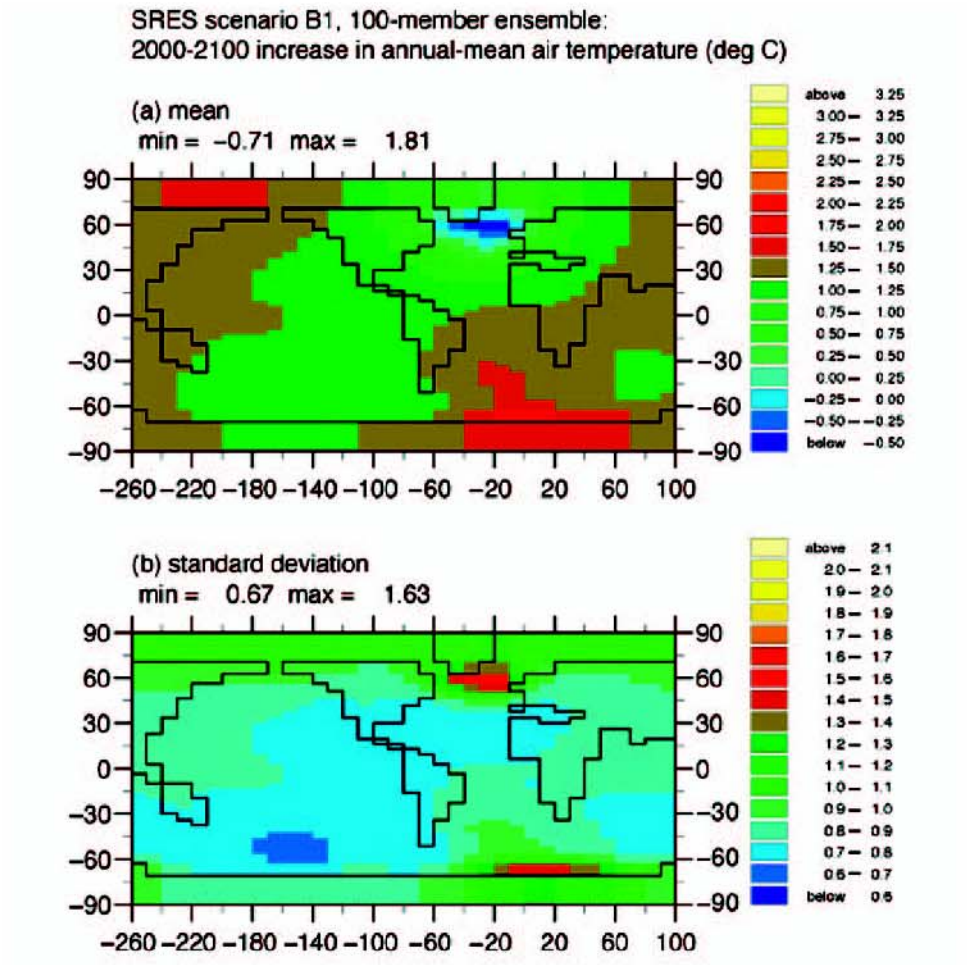

Figure 7.5 Mean and standard deviation of air temperature change in 2100 (relative to 2000) under scenario B1.

change. To illustrate the method we have applied it to an intermediate complexity climate model, C-GOLDSTEIN. The results we obtained were surprising. The probabilities we estimate are much higher than our expectations. *A priori* we expected to obtain probabilities of the order of a few percent or less. The probabilities in Table 7.2 are order 30–40%. There are a number of possible explanations for these differences. Our statistical methodology may be somewhat flawed, the model we have used could be showing unusual behaviour or our *a priori* ideas (and the current consensus) could be wrong. Let us consider each in turn.

The first possibility is that there is a problem with our statistical methodology. The basic method is sound but in our implementation we have made some assumptions and compromises that may influence our results. For example, we have assumed that the input distributions for our parameters are independent of each other and we have discarded the nine runs where the circulation collapsed during spin up. Both of these decisions could have altered our estimated probabilities of collapse. A more thorough elicitation of the input distributions and better sensitivity analysis will enable us to address the problems of specifying input distributions in future work. Moving on to the nine runs that collapsed during the spin up: from measurements we know that the current strength of the Atlantic overturning circulation is in the range 15–20 Sv. When we performed the spin-up, nine of our runs produced current day climates with the overturning circulation approximately zero. We therefore infer that the parameter values used in these runs are not possible. We simply ignored these runs when we built the emulator. This is not correct. When we perform our Monte Carlo simulation we will still be sampling from these regions with parameter sets that we know do not generate the present day climate. Because we discarded those runs, the emulator will interpolate across this region from adjacent parts of parameter space. It is likely that these will themselves have collapsed in 2100 so we may well be overestimating the probability of collapse by including this region. A better procedure would be to build an emulator for the present day and to map out those parts of

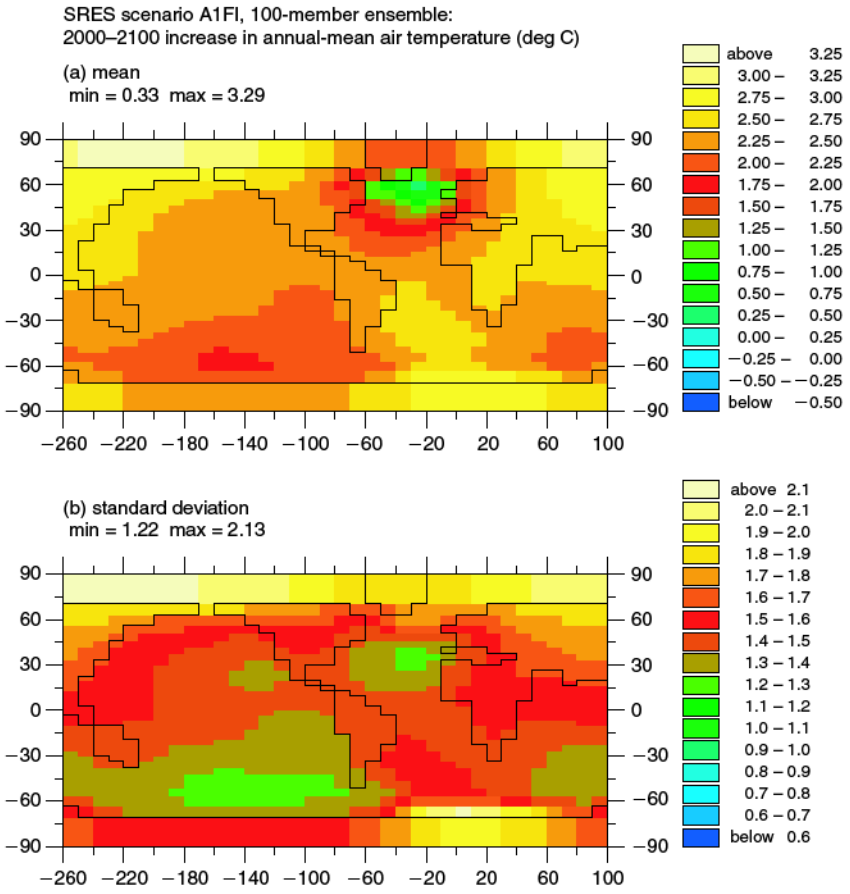

Figure 7.6 As Figure 7.5, under scenario A1FI.

parameter space that result in a collapsed present day circulation. This region could then be set to have zero probability in the input distribution before carrying out the Monte Carlo simulations. This discussion leads us to consider more widely how we might include data in our procedure. The methodology for doing this is explained in Kennedy and O'Hagan (2001).

The second possibility is that the circulation in C-GOLDSTEIN is much more prone to collapse than reality. An intermediate complexity model must by necessity include many assumptions and compromises. A consensus view is that, compared to AOGCMs, the overturning circulation in such models is generally considered more prone to the collapse. However, no one has yet managed to fully explore the behaviour of the overturning circulation across the parameter space of an AOGCM. As discussed above, the spread of our ensemble is not dissimilar to the variation across the set of AOGCMs used by the IPCC. This gives us some confidence that the response of C-GOLDSTEIN's overturning is not very different from the AOGCMs.

The final possibility is that the current consensus is wrong and that the probability of a collapse in the overturning circulation is much higher than believed. There has been little previous work attempting to quantify the probability. Schaeffer et al. (2002) using ECBilt-CLIO, a different intermediate complexity model, state that 'for a high IPCC non-mitigation emission scenario the transition has a high probability', but they do not quantify what they mean by 'high'. Most model runs investigating the collapse of the overturning circulation, such as CMIP, are run at the most likely value for the parameters and therefore approximately at the 50% probability level so would not detect probabilities of collapse of less than 50%. We should, therefore, at least consider the possibility that the current consensus is wrong and that the probability of a shutdown in the overturning circulation is higher than presently believed. However, the most likely reason for our high probabilities is the model we have used is too simple and has omitted important aspects of the climate system. We caution against giving our results too much credence at this stage. However, we believe that our results do show that it is important that quantitative estimates of dangerous, even if unlikely, climate changes can be made. Our calculations need to be repeated with other models and in particular our statistical methodology needs to be extended to make it viable for use with AOGCMs.

Table 7.2 Probability of Atlantic overturning falling below 5 Sv by 2100.

	SRES scenario					
Uncertainty Case	A1B	A2	B1	B2	A1FI	A1T
default uncertainty						
Case 1a	0.37	0.38	0.31	0.32	0.43	0.32
Case 1b	0.38	0.40	0.30	0.31	0.46	0.31
doubled uncertainty in climate sensitivity						
Case 2a	0.37	0.38	0.33	0.33	0.43	0.33
Case 2b	0.39	0.40	0.31	0.32	0.46	0.32
doubled uncertainty in Atlantic-Pacific moisture flux						
Case 3a	0.37	0.38	0.32	0.33	0.43	0.33
Case 3b	0.40	0.40	0.30	0.30	0.46	0.32
doubled uncertainty in CO_2 uptake						
Case 4a	0.38	0.38	0.31	0.32	0.44	0.33
Case 4b	0.38	0.39	0.31	0.31	0.44	0.32
doubled uncertainty in Greenland melt rate						
Case 5a	0.37	0.38	0.31	0.32	0.43	0.32
Case 5b	0.38	0.39	0.30	0.32	0.45	0.32

In Case 1, 'default uncertainty' refers to the standard deviations for all 17 parameters in Table 7.1. In Cases 2–5, 'doubled uncertainty' refers to twice the standard deviation on an individual parameter (italics in Table 7.1). In each case, 'a' ('b') indicates low (high) mean Greenland melt rate.

Acknowledgements

We thank Jonathan Rougier and Tony O'Hagan for discussions and two anonymous referees for their helpful comments. This work was supported by the 'RAPID' directed research programme of the UK Natural Environment Research Council, and by the Tyndall Centre for Climate Change Research.

REFERENCES

Allen, M.R. and D.A. Stainforth. (2002) Towards objective probabilistic climate forecasting. *Nature,* **419**, 228.

Bøggild, C.E., Mayer, C., Podlech, S., Taurisano, A., and Nielsen, S., 2004. Towards an assessment of the balance state of the Greenland Ice Sheet. *Geological Survey of Denmark and Greenland Bulletin* **4**, 81–84.

P.S. Craig, Goldstein, M., Rougier, J.C. and Seheult, A.H., 2001. Bayesian forecasting for complex systems using computer simulators, *Journal of the American Statistical Association*, **96**, 717–729.

Cubasch, U., Meehl, G.A. and 39 others, 2001. Projections of future climate change. *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, J. T. Houghton, et al. Eds., Cambridge University Press, Cambridge, UK, 525–582.

Edwards, N.R. and Marsh, R., 2005. Uncertainties due to transport-parameter sensitivity in an efficient 3-D ocean-climate model. *Clim. Dyn.,* **24**, 415–433.

Garthwaite, P.H., Kadane, J.B. and O'Hagan, A., 2005. Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association,* **100**, 680–701.

Hargreaves, J.C., Annan, J.D., Edwards, N.R. and Marsh, R., 2004. An efficient climate forecasting method using an intermediate complexity Earth system model and the ensemble Kalman Filter. *Clim. Dyn.,* **23**, 745–760.

Huybrechts, P. and de Wolde, J., 1999. The dynamic response of the Greenland and Antarctic ice sheets to multiple-century climatic warming. *J. Climate,* **12**, 2169–2188.

Johnston, W.R., 2004. Historical data relating to global climate change. Available from http://www.johnstonsarchive.net/environment/co2table.html

Kennedy, M.C. and O'Hagan, A., 2001. Bayesian calibration of computer models *Journal of the Royal Statistical Society Series B – Statistical Methodology,* **63**, 425–450.

Knutti, R., Stocker, T.F., Joos, F. and Plattner, G.-K., 2003. Probabilistic climate change projections using neural networks. *Clim. Dyn.,* **21**, 257–272.

McKay, M.D., Beckman, R.J., and Conover, W.J., 1979. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics,* **21**, 239–245.

Marsh, R., Yool, A., Lenton, T.M., Gulamali, M.Y., Edwards, N.R., Shepherd, J.G., Krznaric, M., Newhouse, S. and Cox, S.J., 2004. Bistability of the thermohaline circulation identified through comprehensive 2-parameter sweeps of an efficient climate model. *Clim. Dyn.,* **23**, 761–777.

Marsh, R., De Cuevas, B.A., Coward, A.C., Bryden, H.L and Alvarez, M., 2005. Thermohaline circulation at three key sections in the North Atlantic over 1985–2002. *Geophys. Res. Lett.,* **32**, doi:10.1029/2004GL022281.

Marotzke, J., 2000. Abrupt climate change and thermohaline circulation: Mechanisms and predictability. *Proceedings of the National Academy of Sciences of the United States of America,* **97**, 1347–1350.

Nakicenovic, N. and Swart, R. (eds.), 2000. *Special Report on Emissions Scenarios*. Cambridge University Press, Cambridge, UK, 612 pp.

Oakley, J. and O'Hagan, A., 2002. A Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika,* **89**, 769–784.

Prentice, I.C. and 60 others, 2001. The Carbon Cycle and Atmospheric Carbon Dioxide. *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change,* Houghton, J.T. et al. Eds., Cambridge University Press, Cambridge, UK, 183–237.

Rahmstorf, S. and Ganopolski, A., 1999. Long-term global warming scenarios computed with an efficient coupled climate model. *Clim. Change,* **43**, 353–367.

Rind, D., Peteet, D., Broecker, W., McIntyre, A. and Ruddiman, W., 1986. The impact of cold North Atlantic sea surface temperatures on climate: Implications for the Younger Dryas cooling (11–10 k). *Clim. Dyn.,* **1**, 3–33.

Stainforth, D.A., Aina, T., Christensen, C. and 13 others, 2005. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, **433**, 403–406.

Schaeffer, M., Selten, F.M., Opsteegh, J.D. and Goosse, H., 2002. Intrinsic limits to predictability of abrupt regional climate change in IPCC SRES scenarios. *Geophys. Res. Lett.,* **29**, doi:10.1029/2002GL015254.

Vellinga, M. and Wood, R.A., 2002. Global climate impacts of a collapse of the Atlantic thermohaline circulation. *Climatic Change,* **54**, 251–267.

Wigley, T.M.L. and Raper, S.C.B., 2001. Interpretation of high projections for global-mean warming, *Science,* **293**, 451–454.