

# An efficient climate forecasting method using an intermediate complexity Earth System Model and the Ensemble Kalman Filter

J.C. Hargreaves<sup>1\*</sup>, J.D. Annan<sup>1</sup>, N. R. Edwards<sup>2</sup>, R. Marsh<sup>3</sup>

<sup>1</sup> Frontier Research System for Global Change  
3173-25 Showa-machi, Kanazawa-ku  
Yokohama, Kanagawa 236-0001, Japan  
and

Proudman Oceanographic Laboratory, Joseph Proudman Building  
6 Brownlow Street, Liverpool, L3 5DA, UK

<sup>2</sup> Climate and Environmental Physics, Physics Institute,  
University of Bern, Sidlerstrasse 5, CH-3012 Bern, Switzerland

<sup>3</sup> James Rennell Division, Southampton Oceanography Centre,  
Empress Dock, Southampton SO14 3ZH, UK

\* Corresponding Author  
email: [jules@jamstec.go.jp](mailto:jules@jamstec.go.jp)  
fax: +81-45-778-5707

20 July 2004

## Abstract

We present the implementation and results of a model tuning and ensemble forecasting experiment using an ensemble Kalman filter for the simultaneous estimation of 12 parameters in a low resolution coupled atmosphere-ocean Earth System Model by tuning it to realistic data sets consisting of Levitus ocean temperature/salinity climatology, and NCEP/NCAR atmospheric temperature/humidity reanalysis data. The resulting ensemble of tuned model states is validated by comparing various diagnostics, such as mass and heat transports, to observational estimates and other model results. We show that this ensemble has a very reasonable climatology, with the 3-D ocean in particular having comparable realism to much more expensive coupled numerical models, at least in respect of these averaged indicators.

A simple global warming experiment is performed to investigate the response and predictability of the climate to a change in radiative forcing, due to 100 years of 1% per annum atmospheric CO<sub>2</sub> increase. The equilibrium surface air temperature rise for this CO<sub>2</sub> increase is  $4.2 \pm 0.1^\circ\text{C}$ , which is approached on a time scale of 1000 years. The simple atmosphere in this version of the model is missing several factors which, if included, would substantially increase the uncertainty of this estimate. However, even within this ensemble, there is substantial regional variability due to the possibility of collapse of the North Atlantic thermohaline circulation (THC), which switches off in more than one third of the ensemble members. For these cases, the regional temperature is not only 3–5°C colder than in the warmed worlds where the THC remains switched on, but is also 1–2°C colder than the current climate.

Our results, which illustrate how objective probabilistic projections of future climate change can be efficiently generated, indicate a substantial uncertainty in the long-term future of the THC, and therefore the regional climate of western Europe. However, this uncertainty is only apparent in long-term integrations, with the initial transient response being similar across the entire ensemble. Application of this ensemble Kalman filtering technique to more complete climate models would improve the objectivity of probabilistic forecasts and hence should lead to significantly increased understanding of the uncertainty of our future climate.

# 1 Introduction

In contrast to short-term operational weather prediction, climate forecasts (for a given set of boundary conditions, ie a specific scenario) depend more strongly on parameterisations and less strongly on initial conditions. At the multi-decadal time scale and beyond, model estimates of the response to anthropogenic forcing (as measured by for example the typical indicator of global mean surface air temperature) substantially exceed the range of natural variability (Collins & Allen, 2002). However, the magnitude of the response depends greatly on the nature of the parameterisations contained within the model, and the particular parameter values selected (Houghton *et al.*, 2001). Therefore, in order to generate meaningful predictions, it is important that the parameter values should be tuned to appropriate values. Moreover, given that no deterministic prediction will ever be exactly correct, it is also important to quantify the uncertainty associated with a forecast, which itself depends on the confidence with which parameter values can be determined. For these reasons, the problem of parameter estimation in climate modelling has recently attracted a great deal of attention (eg Forest *et al.*, 2000; Andronova & Schlesinger, 2001; Knutti *et al.*, 2002; Gregory *et al.*, 2002). We have introduced, in Annan *et al.* (2004) and Annan & Hargreaves (2004), a new parameter estimation system based on the ensemble Kalman filter (EnKF) (Evensen, 1994; Keppenne, 2000). The potential of the system was clearly demonstrated by application to identical twin testing with a new Earth system model of intermediate complexity (EMIC). We were able to simultaneously estimate 12 parameters from observations of the climatological mean model state, with an ensemble of only 54 model runs. This represents a large increase in efficiency when compared to the rather simple Bayesian and brute-force sampling methods that have been previously used. Identical twin testing, where the model itself is used to generate synthetic data, is a useful step in the development of assimilation methods, but can be a rather weak test. In that type of experiment, the error statistics of the surrogate observations are known precisely, and the perfect model assumption is not challenged. Both of these factors may play an important role in real applications using observational data. In this paper, we demonstrate the application of the parameter estimation method in tuning the EMIC to realistic climatology consisting of ocean temperature and salinity, from 1945 to 1998 (Levitus, 1998), and NCEP/NCAR atmospheric reanalysis data (surface air temperature and humidity only) averaged between 1948-2002 (NCEP Reanalysis data provided by the NOAA-CIRES Climate Diagnostics Center, Boulder, Colorado, USA, from their Web site at <http://www.cdc.noaa.gov/>).

We briefly introduce the model, data and method in Section 2. The importance of model error is discussed, and we describe a simple method to account for it. In Section 3, we present the results of tuning the model to the climatological data, and validate it by comparison with various observational and model-based

estimates of heat and mass transports. Although this model has a rather low resolution and a highly simplified representation of many physical processes, it appears to give a realistic description of large-scale circulation, comparable to the much more expensive and complex models of CMIP (Coupled Model Intercomparison Project, <http://www-pcmdi.llnl.gov/cmip/>). This in itself is a clear illustration of the power of the tuning procedure.

We further investigate the response of the ensemble to some idealised anthropogenic CO<sub>2</sub> emissions scenarios in Section 4. The effect of changing atmospheric CO<sub>2</sub> is parameterised as a direct radiative forcing of 4Wm<sup>-2</sup> for each doubling of CO<sub>2</sub>, and the model omits a large variety of carbon-cycle and atmospheric dynamical feedbacks such as the highly uncertain aerosol and cloud feedback mechanisms, so this experiment is best considered as a demonstration of the potential of the method rather than a definitive forecast of future climate change. Even though the globally averaged climate change is highly constrained due to the simplicity of the model, there is significant local uncertainty due primarily to the possible collapse of the North Atlantic thermohaline circulation (THC). The probability of a collapse of the North Atlantic THC is quantified, with the outcome highly sensitive to the level of CO<sub>2</sub> reached. We end with some conclusions in Section 5.

## 2 Model, data and method

### 2.1 Model

The model, C-GOLDSTEIN, was outlined in Annan *et al.* (2004) and is described more fully in Edwards & Marsh (2003). Briefly, it is a prototype EMIC being constructed as part of the GENIE project (<http://www.genie.ac.uk/>). The version used here is a coupled atmosphere-ocean general circulation model on a 36 × 36 equal-area horizontal grid. The resolution and topography along with some model output is shown in Figures 1 and 2. The ocean is a 3-D frictional geostrophic model on an 8-level  $z$ -coordinate grid (Edwards & Shepherd, 2002) and the atmosphere is a 2-D energy/moisture balance model similar to that of Weaver *et al.* (2001). The model also contains a simple thermodynamic and dynamic sea ice formulation following Semtner (1976) and Hibler (1979). Further components including biogeochemical cycles, dynamical land ice sheets and a low-resolution 3-D dynamical atmosphere are currently under development. A major feature of this model is its computational efficiency, which allows integration to equilibrium ( $\sim 2,000$  model years from a cold start) in only a few hours on a desktop computer.

We tune the same set of 12 parameters as in Edwards & Marsh (2003) and Annan *et al.* (2004), and for clarity we repeat their description here. Ocean temperature

and salinity are diffused along and across isopycnal surfaces with constant diffusivities, and advected by the frictional geostrophic velocity field. The frictional drag coefficient (representing nonlinear momentum terms) is spatially variable, increasing near steep topographic features and near the equator from a constant interior value. To counteract the dissipative effect of the drag and the uncertainty of the momentum flux from the atmosphere to the ocean (which depends on the surface roughness) we allow for a scaling factor of the wind stress. In the one-layer atmosphere, diffusivity of surface specific humidity is constant, while diffusivity of surface temperature is an asymmetric function of latitude. Here we wish to consider the effect of varying the width and asymmetry of this function and thus we define the temperature diffusivity  $\kappa_T$  via

$$\kappa_T = k_T \left( s_d \frac{2\theta + \pi}{\pi} + (1 - s_d) \frac{\exp(-(\theta/l_d)^2) - c}{1 - c} \right), \quad (1)$$

where  $\theta$  is latitude (in radians),  $l_d$  and  $s_d$  are width and slope parameters and the constant  $c$  is given by

$$c = \exp \left( - \left( \frac{\pi}{2l_d} \right)^2 \right). \quad (2)$$

The vertically averaged effect of atmospheric advective transport is represented by advection with a fixed surface wind field derived from NCEP analysis, scaled by a coefficient, following Weaver *et al.* (2001), which takes different values for moisture advection and for the zonal advection of heat. There is no meridional advection of heat in the model atmosphere. To compensate for a lack of inter-basin moisture transport produced by our simplified atmosphere we introduce a constant redistribution of moisture from the surface Atlantic to Pacific following the pattern observed by Oort (1983). This ‘flux adjustment’ is an important parameter of the model hydrological cycle, but is unrelated to the flux adjustments required in early coupled models to prevent climate drift, which is not a problem for this model. Sea ice height and fractional area are advected by the surface ocean currents and diffused with constant diffusivity.

With two diffusivities in the ocean and one for sea ice, an Atlantic-Pacific moisture flux adjustment, two parameters controlling wind-driven circulation and six parameters controlling the atmospheric heat and moisture transport, we have a set of 12 model parameters which are allowed to vary. The prior distributions for all 12 parameters are given in Table 1, along with the results that will be discussed in Section 3. The ranges of the distributions are chosen to cover, or exceed, a range of reasonable choices of appropriate values for such a model as discussed by Edwards & Marsh (2003). When the model is repeatedly run using parameters chosen independently at random from these priors, the climatologies produced span a large range which is generally well in excess of the uncertainty of the true climate state (some diagnostics of this prior are mentioned in Section 3).

Thus, we are deliberately starting from a position of substantial ignorance in order to investigate what can be objectively determined from the data alone.

## 2.2 Method

The assimilation method is an iterative application of the EnKF, and is described more fully in Annan *et al.* (2004). The climatological parameter estimation problem studied in this paper is a steady state problem, somewhat different in detail (and in principle simpler) than the more conventional time-varying implementations of the EnKF. However, the prior assumption of substantial ignorance, combined with the nonlinearity of the model and high dimensionality of the parameter space being explored, means that a direct solution of the steady state problem does not work well. Therefore, an iterative scheme has been developed which repeats a cycle of ensemble inflation (in which the spread of the ensemble is increased by increasing the distance from each member to the ensemble mean by a fixed multiplicative factor), data assimilation and model integration over a specified time interval, in order to converge to the final solution. As demonstrated in Annan *et al.* (2004) and Annan & Hargreaves (2004), this iterative method converges robustly to the correct solution in identical twin testing. All of the experiments described here used an assimilation cycle of 100 years length, with an expansion factor of 5%, and an ensemble size of 54, all of these values (which do not significantly affect the converged solution) being chosen primarily for computational convenience.

Using only steady state data eliminates any information concerning the rate of change that may be contained in recent time series of observational data. This is clearly undesirable, especially since the ultimate purpose of the tuning is to estimate the rate of future climate change. In principle, the assimilation method used here can also be applied to time series data, since the augmented model state method allows for arbitrary asynoptic observations to be used (note that climatological observations are themselves asynoptic). Further development of the method to this end is under way.

There is one important detail in the assimilation that should be explained here. The previous work considered only an identical twin experiment in which the ‘strong constraint’ assumption of a perfect model is not violated. In any real application, the model will contain some inaccuracies and approximations over and above those relating to imperfect parameter values; for example numerical diffusion, truncation errors, and, more generally, many approximations to the processes being modelled. This ‘model error’ formally violates the assumptions underlying many analysis and estimation algorithms (including this implementation of the EnKF), and on a practical level, if the model error is significant, then this can degrade the quality and reliability of the results. In particular, the results

will tend to have an unrealistically narrow uncertainty associated with them — the problem of ‘false confidence’ which was illustrated in Hargreaves & Annan (2002). In that work, the true forecast (and even hindcast) error was substantially greater than the width of the ensemble generated by an objective estimation scheme. A simple diagnostic for the occurrence of this problem is to examine the residuals after fitting the model. If they are larger than expected (as defined by a high  $\chi$ -square value (Press *et al.*, 1994, 15.2)) or are significantly correlated in time and/or space, then these are indications that the model contains significant structural errors that should be accounted for. Thacker (2003) discusses the detection of model-data incompatibility in more detail. One possible approach to account for this problem is to consider the discrepancy between model and data as ‘representation error’: that is, the model does not represent the real system perfectly, but rather is a somewhat erroneous approximation to it. In this case, the model state does not represent an estimate of the truth but rather can only be considered to be an estimate of the projection of truth onto the manifold spanned by the model. The correct treatment of representation error (where it cannot be filtered from the data, as in for example de-tiding of sea surface elevation data) is to add it to the estimated observation error of the data (Fukumori, 2001), thereby effectively decreasing the assumed accuracy of the observations and consequently increasing the width of the estimated pdf. Since the magnitude and correlation length scale of the representation error are not known *a priori* but can only be diagnosed through an *a posteriori* examination of the differences between the model and data, this remains a somewhat tunable factor in this application, and we have validated our choice by comparing the width of the resulting ensemble to the uncertainty of various observational estimates in Section 3. A largely equivalent technique to limit data influence that is sometimes adopted in assimilation studies is to use a temporally and/or geographically sparse subset of the data, but here the model resolution is so low that we do not wish to artificially generate any data-free regions or other artefacts due to subsampling.

## 2.3 Data

Figures 1 and 2 show some of the data, averaged onto the model grid. The ocean data set used consists of fully three-dimensional temperature and salinity fields, but only the top level is shown here. The atmospheric data are two-dimensional (surface air temperature and humidity). Both data sets represent an average over roughly five decades with a near-complete overlap. During this time, there has been significant anthropogenic perturbation to the atmospheric CO<sub>2</sub> concentration. Since the response of the atmosphere and upper ocean is much more rapid than that of the deep ocean, these data sets do not precisely represent a steady-state climatology appropriate to any fixed CO<sub>2</sub> level. However, the mean overall perturbation from the pre-industrial state is certainly less than 0.5°C (this being a

typical estimate for the recent anthropogenically-forced surface warming (Hansen *et al.*, 2002), with the temperature rise being much lower in the ocean interior (Levitus *et al.*, 2000)), and since this discrepancy is small compared to other faults in our model, we discount this minor problem for the purposes of this study. The time-varying assimilation procedure which is being developed will enable us to address this issue with more precision.

These data are by no means the only observations that could help us to determine the full model state and dynamics (and therefore parameter values). For example, sea ice, precipitation, heat and mass transports and passive tracer distributions could also all be helpful in principle. However, the data selected do provide a substantial volume of reliable information over the entire global domain, and have been widely used and well validated. As such, they are a suitable basis for the preliminary experiment performed here. Since we do not know the covariances of the errors on the data, we assume that the observational errors on the different data types are proportional to the square root of their variances as in Annan *et al.* (2004) and ignore spatial correlations. To account for the model error problem mentioned in Section 2.2, we fix an observational error of 3°C for the ocean data, and scale the error on the other data types appropriately. This value is substantially larger than the true observational uncertainty, but this deliberate overestimate compensates for the representation errors in the model, and results in an ensemble width that appears to plausibly represent the true uncertainty of the climate system state, as we now show in Section 3.

## 3 Results and Validation

### 3.1 Convergence

The mean atmospheric and ocean surface fields from the converged ensemble are shown in Figures 1 and 2. These results were obtained after 150 iterations (representing 15,000 years of integration for each ensemble member). However, there was very little change beyond 20 iterations (2,000 years, comparable to the equilibration time scale of the deep ocean). The rate of convergence of model variables is similar to that seen in the identical twin test (not shown here). For application to more expensive models, this process could be substantially speeded up by initialising from a realistic climatology (rather than the uniform fields used here) and by allowing for a modest drift in the deep ocean climate, which is not necessarily in true equilibrium in any case and will not significantly affect  $O(100)$  year forecasts.

As a single scalar measure of the goodness of fit, we use the cost function described in Edwards & Marsh (2003). This is a global root mean square error, with each

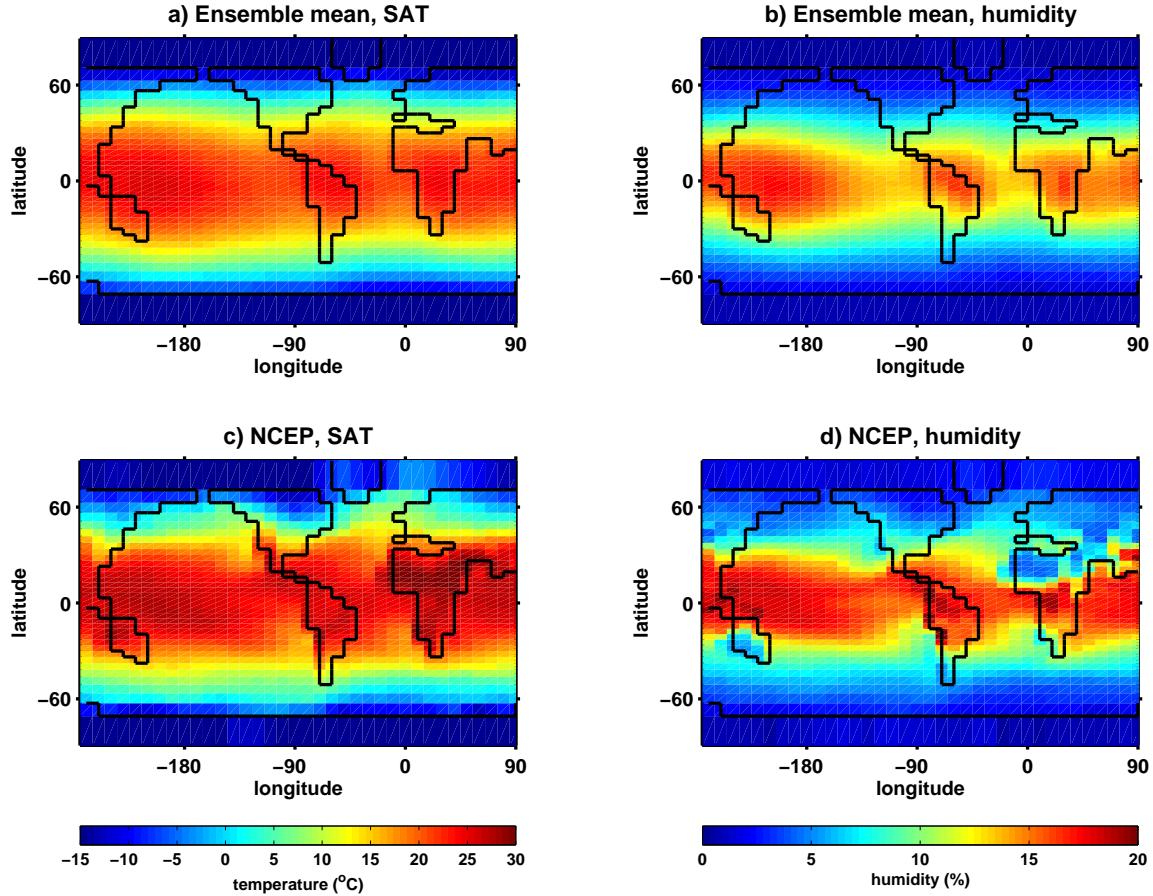


Figure 1: NCEP reanalysis surface air temperature (SAT) and humidity, with corresponding ensemble mean output

data type normalised by the square root of its spatial variance. Therefore, spatially uniform model fields with the correct global mean would have a cost of 1, and our final figure of around 0.5 (Figure 3) indicates that typically three-quarters of the variance of each variable type is explained by the model. This compares favourably with the results from the Latin hypercube ensemble experiment of Edwards & Marsh (2003), in which the costs of 1000 samples ranged from 0.61 to 2.8. As can be seen from Figure 3, the cost converges to its final value in around 20 iterations.

### 3.2 Parameter values

Figure 4 shows the evolution of the mean (central lines) and one standard deviation width (outer lines) of the parameter values. The horizontal lines at time  $t < 0$  in each plot indicate the prior distributions, selected to cover a wide range of plausible values. Convergence here is slightly slower than examination of the cost

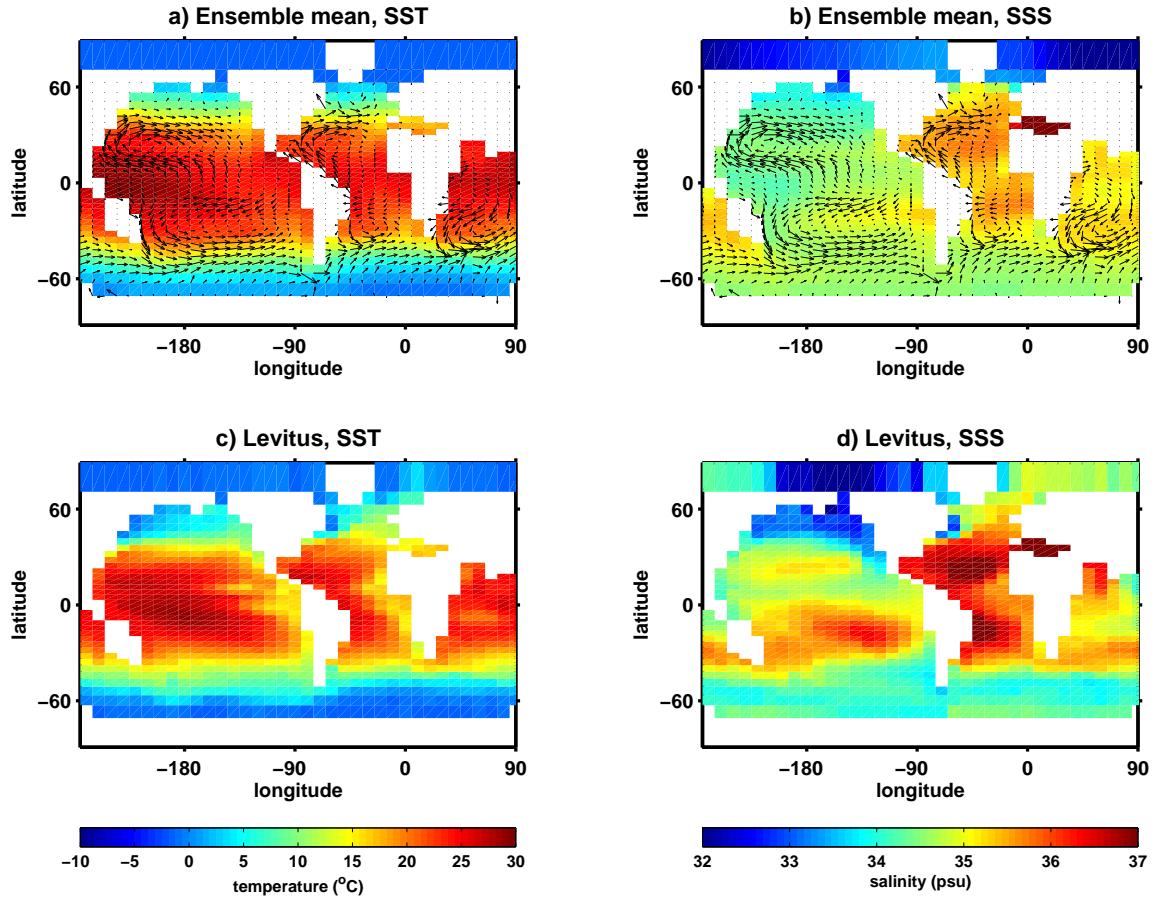


Figure 2: Levitus sea surface temperature (SST) and salinity (SSS), with corresponding ensemble mean output and surface currents

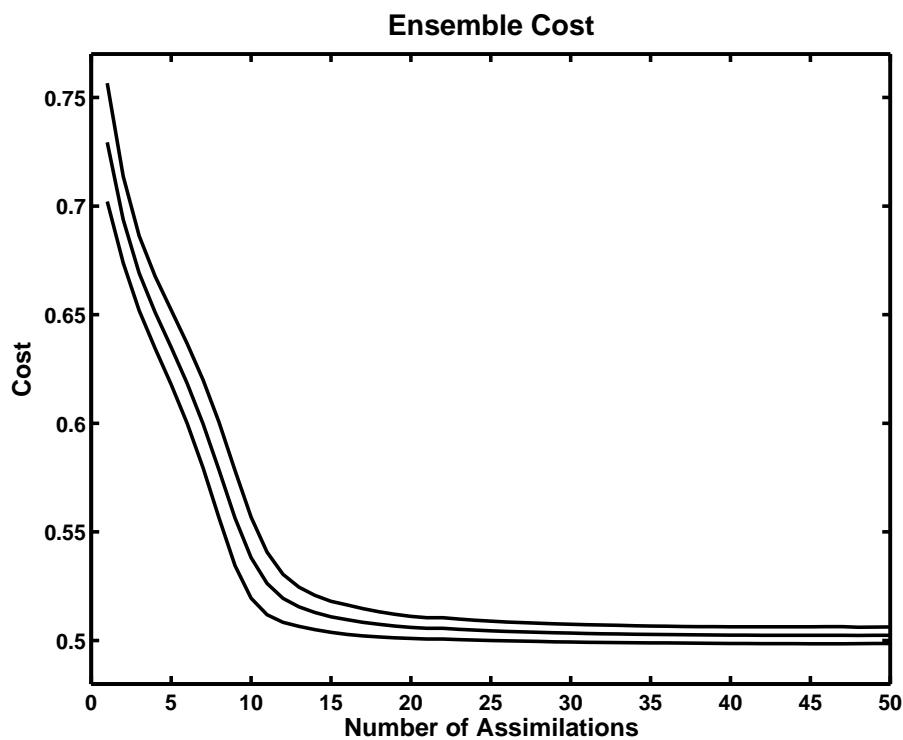


Figure 3: Decrease in cost function (a global measure of model-data error), ensemble mean and one standard deviation range.

function alone would indicate, but is still close to completion after 50 cycles. The initial spikes in several of the parameter values are due to the unrealistic and unsteady model state at this time.

Parameter	Prior		Posterior	
	Mean	std.dev.	Mean	std.dev.
<b>Ocean</b>				
Wind-scale	1.9	0.4	1.7	0.2
Isopycnal diffusion ( $\text{m}^2\text{s}^{-1}$ )	5100	2300	4100	800
$\log_{10}(\text{diapycnal diffusion } (\text{m}^2\text{s}^{-1})/2 \times 10^{-5})$	0.07	0.44	-0.04	0.27
1/friction (days)	2.7	0.9	3.4	0.7
<b>Atmosphere</b>				
T diffusion amplitude ( $\text{m}^2\text{s}^{-1}$ )/ $10^6$	6.9	1.5	3.8	0.5
T diffusion width (radians)	1.3	0.78	1.3	0.2
$\log_{10}(\text{T diffusion slope}/0.1)$	-0.32	0.36	-0.16	0.30
T advection coefficient	0.35	0.17	0.06	0.05
$\log_{10}(\text{Q diffusion } (\text{m}^2\text{s}^{-1})/10^5)$	0.57	0.43	1.2	0.04
Q advection coefficient	0.40	0.19	0.14	0.04
FWF adj (Sv)	0.32	0.16	0.29	0.03
<b>Ice</b>				
Sea ice diffusion ( $\text{m}^2\text{s}^{-1}$ )	3900	1800	6200	1500

Table 1: Prior and posterior distributions of the parameters

As in the identical twin testing described by Annan *et al.* (2004), some parameters are hardly constrained by the data, that is to say, the posterior distributions are not much narrower than the priors. For example, the ice diffusion coefficient is very uncertain. Even though the ensemble has an acceptable representation of ice (described below in Section 3.6), this is essentially determined by the local sea surface and atmospheric temperatures, and its dynamics play a minor rôle. At the other extreme, it is reassuring that the freshwater flux adjustment (a redistribution of water from the Atlantic to the Pacific basins, to correct for our atmosphere's inability to perform this function) is constrained so as to give a total moisture flux (when added to the model's own modest contribution of around 0.02Sv) very close to the observational estimate of 0.32Sv (Oort, 1983).

There are a handful of significant correlations between the parameter values according to the sample statistics of the ensemble, similar but not identical to those found in the previous identical twin test. The width and amplitude of the atmospheric temperature diffusion coefficients have a significant negative correlation

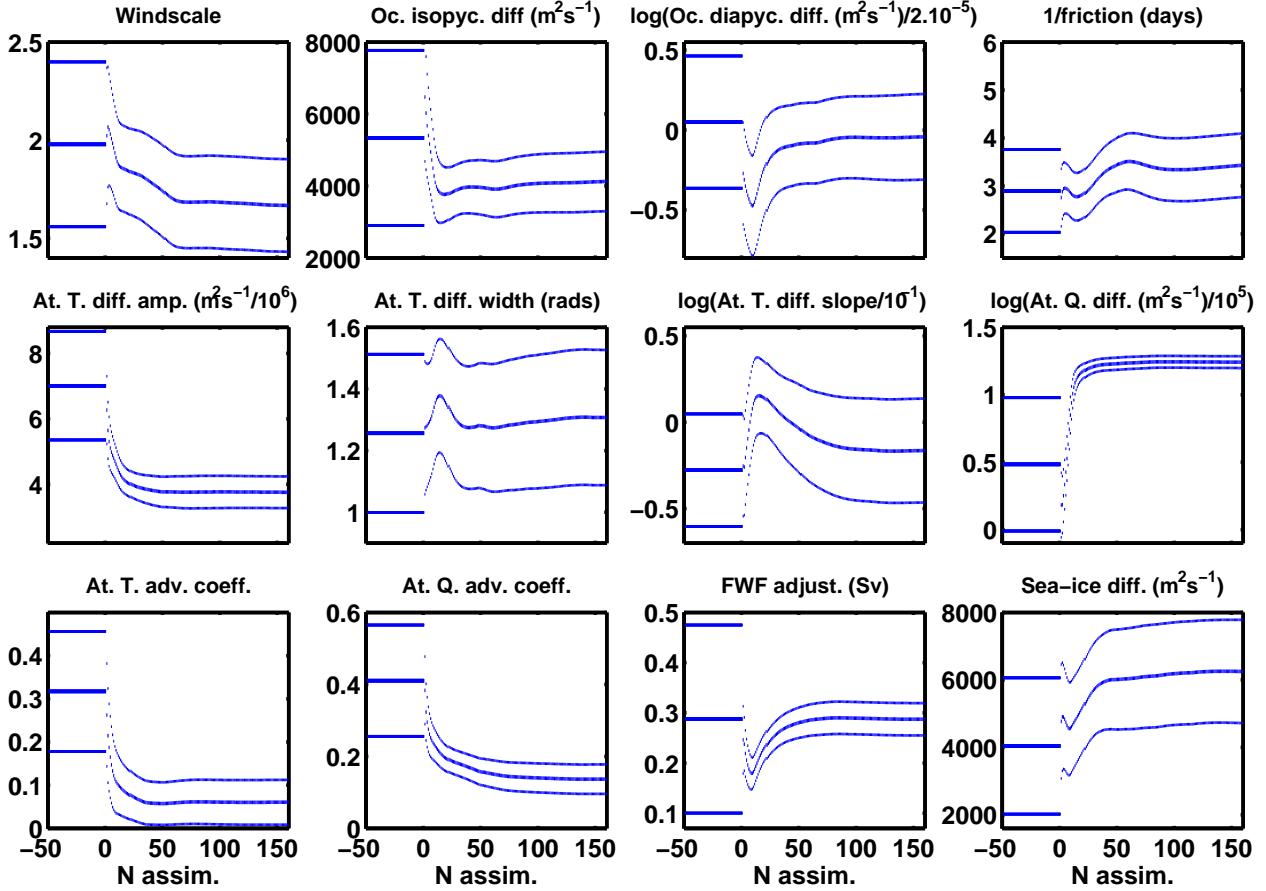


Figure 4: Evolution of parameter values, ensemble mean (central lines) and one standard deviation widths ( $\log = \log_{10}$  throughout).

( $r^2 = -0.6$ ), as might be expected from their roughly equivalent effects on large-scale heat transport. The ocean inverse friction coefficient, and wind scaling factor, have a negative correlation of -0.4, which is again unsurprising as they have a largely equivalent action on ocean circulation. In the identical twin test, this negative correlation was also present but did not meet the threshold for statistical significance which is an absolute magnitude of 0.27 or greater (for 90% confidence). The ocean isopycnal diffusion coefficient, atmospheric moisture diffusion coefficient and fresh water flux correction are all positively correlated with one another, with values ranging from 0.5 to 0.7. These results are slightly different from those of the identical twin test, but the correlated parameters are all strongly implicated in determining the strength of the meridional overturning circulation and thus have somewhat interchangeable effects.

### 3.3 State variables

The global root mean square (RMS) errors of the ensemble mean of each field, in physical units, are  $1.45^{\circ}\text{C}$  (ocean temperature),  $0.22\text{psu}$  (ocean salinity),  $1.31^{\circ}\text{C}$  (atmospheric temperature) and  $0.0011$  (atmospheric humidity, dimensionless mass ratio of water to air). After normalisation by the square root of the variance, the relative salinity error is greater than the other three data types combined, probably due to rather poor moisture transport in the atmosphere which will be discussed in Section 3.5.

The validity of our climate state is investigated by comparison with observational and model estimates from various sources. Some results from the Coupled Model Intercomparison Project (CMIP), as summarised by Lambert & Boer (2001) and Jia (2003), provide a particularly useful benchmark. The CMIP project consists of a wide-ranging investigation into many aspects of the behaviour of many of the most widely used coupled climate models, under both steady conditions and an idealised anthropogenic forcing scenario. In this initial validation, we are only concerned with the model climatologies under steady pre-anthropogenic conditions. Lambert & Boer (2001) summarises the basic climate variables from CMIP1, the first phase of CMIP. Their paper contains comparisons of model fields and data at a limited number of locations. Although global numerical RMS error statistics are not provided, it seems that our ensemble mean compares acceptably well with the CMIP1 models in terms of the basic state variables such as temperature and salinity. For example, the CMIP1 ocean temperature profiles at  $15^{\circ}\text{S}$  (Lambert & Boer, 2001, Figure 6) appear to have a typical error in the region of  $1\text{--}2^{\circ}\text{C}$ , even after zonal averaging which will underestimate the pointwise RMS error. The salinity field in our model is rather poor, and it is clear from Figure 2 that the meridional gradients in the surface salinity fields are substantially weaker than the observations. This is due to the inability of the simple atmosphere to transport moisture effectively from the saline equatorial region to the fresher poles, and is discussed further in Section 3.5. Despite this problem, our global RMS error of  $0.22\text{psu}$  compares very reasonably with the CMIP1 results at the surface and  $1000\text{m}$  depth (Lambert & Boer, 2001, Figures 8 and 9, again the zonal averages plotted will underestimate the pointwise error). The atmospheric surface air temperature plots of Lambert & Boer (2001, Figures 3 and 4) also show comparable errors to our global RMS value.

Since these variables were directly assimilated into our ensemble, it is hardly surprising that the agreement between model and data is reasonable. A more severe and interesting test is to examine derived quantities such as heat and mass transports, since the transports must be generated by the model dynamics and therefore these quantities are not purely dependent on the assimilated data. In the remainder of this section, we discuss the transports in more detail.

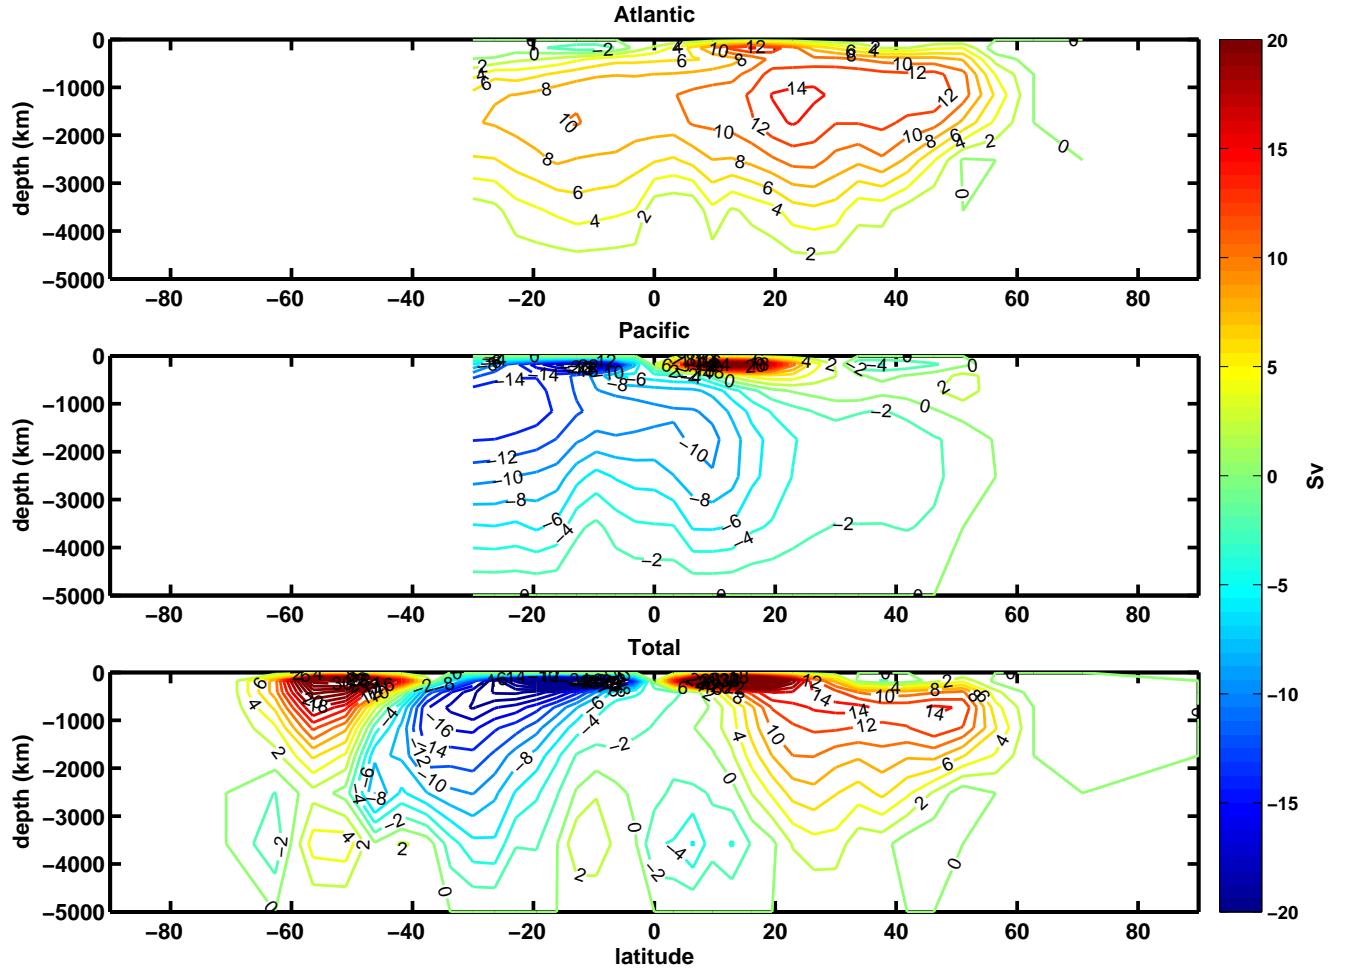


Figure 5: Zonally averaged overturning (Sv)

### 3.4 Ocean circulation and transport

Since the ocean is the most sophisticated component of our coupled model, we focus most of our attention there. Jia (2003) analysed the output from the ocean components in eighteen of the CMIP ‘control’ runs (current climatology), and compared various transports to observational estimates at the two latitudes 25°N and 30°S in the different ocean basins.

The zonally-averaged basin-wide and global overturning of the mean of our ensemble are shown in Figure 5, and the heat transports are plotted in Figure 6. In the following subsections we examine these results in more detail, focussing on each of the major ocean basins in turn.

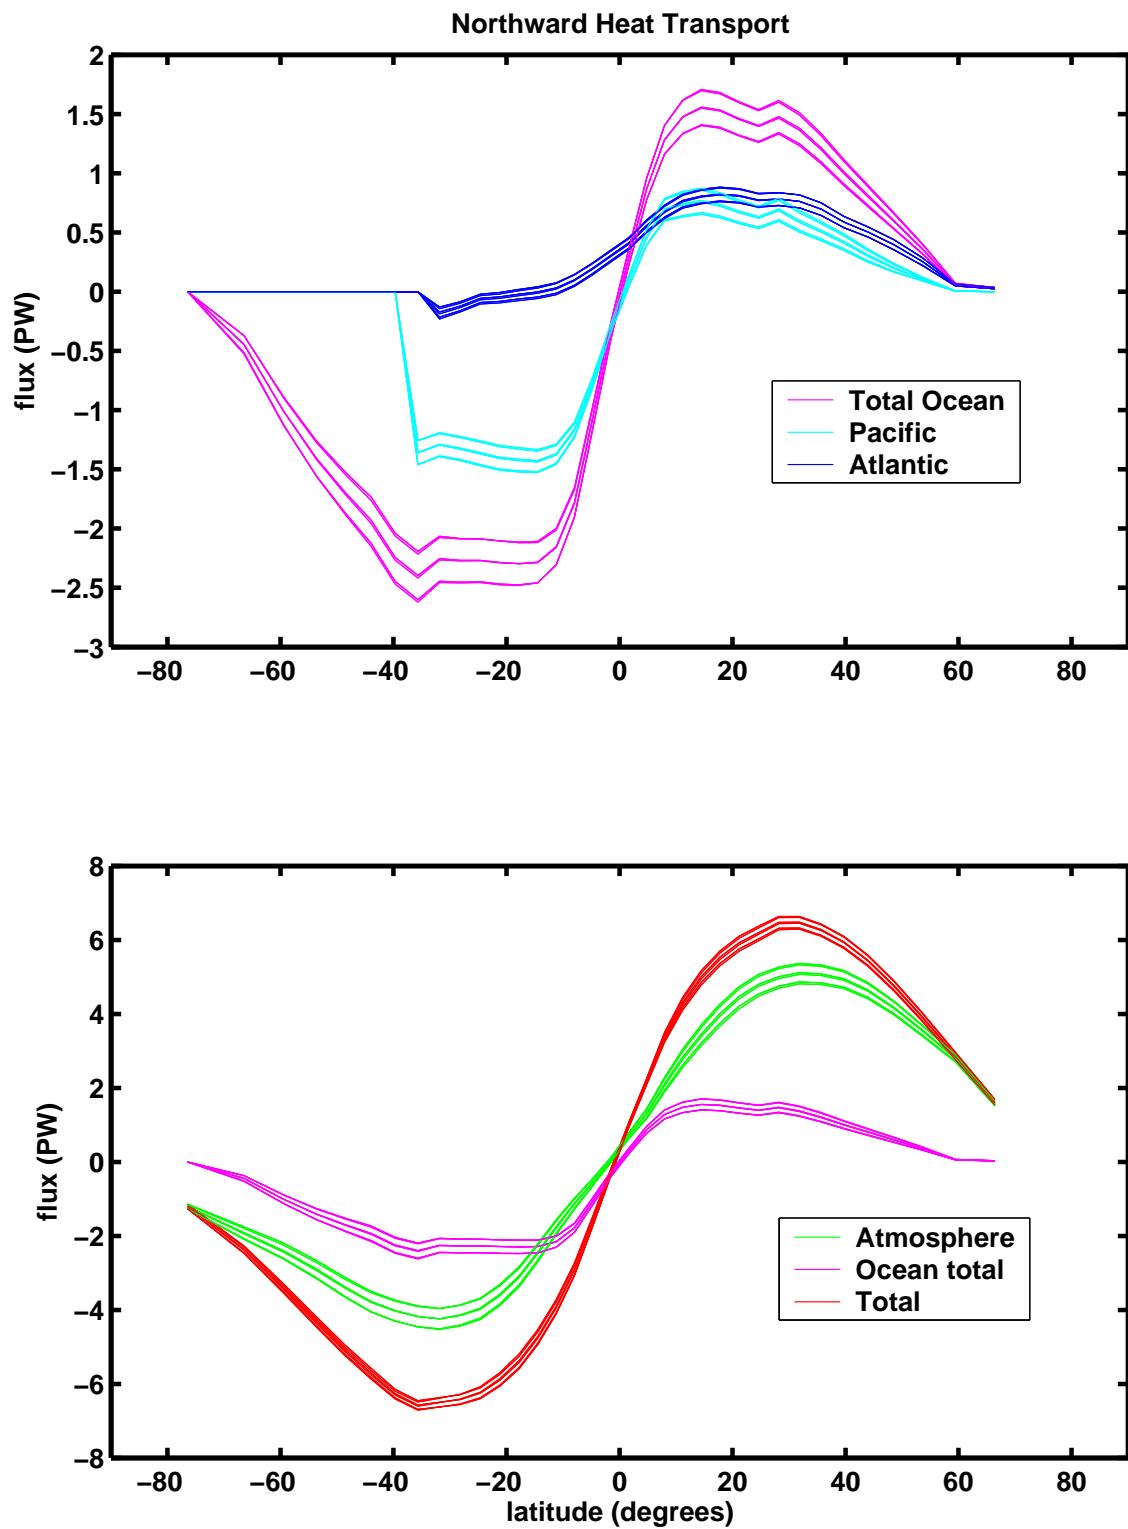


Figure 6: Zonally averaged northward heat transports

### 3.4.1 Atlantic Ocean

Observational estimates of the zonally integrated climatological heat and mass transports at 25°N in the Atlantic ocean are shown as the large red dots in the upper plot of Figure 7, along with the results from the CMIP experiment (black crosses) and our own ensemble members (dark blue crosses). The cyan crosses will be discussed in Section 4. The CMIP and observational data are taken from Jia (2003), from whose Figures 4 and 5 this diagram was derived. Sources and numerical values for the CMIP models and observational data are listed in Jia (2003, Table 2).

The observational estimates of heat transport are around 1.1–1.3PW, with an overturning strength of 16–20Sv. However, few of the CMIP models are close to these values, with only 3 of them exceeding 1PW, and they also have a very wide range of overturning strengths. According to Jia (2003), the basic reason for the unrealistically low heat transports in the CMIP models is the inability of these models to represent the vertical heat gradient correctly, with most being much too warm in the deep ocean (Figure 7, lower plot).

In contrast, our model generates a much more realistic temperature distribution, with the ensemble mean upper and lower temperatures being 18.3 and 3.3°C. However, the overturning of  $14.5 \pm 1.5$ Sv is slightly too small, and this results in a low heat transport of  $0.8 \pm 0.1$ PW. It is straightforward to generate an ensemble which has closer agreement with these, and other, transport estimates (by assimilating the estimates directly into the ensemble), but this would prevent us from using these estimates as independent validation of the model, as well as potentially suffering from the problem of data over-use since the transport estimates are often based to some extent on the climatological fields that we are already assimilating.

The distribution of our ensemble members in the upper plot agrees well with the estimate from Böning *et al.* (1996) that for each 2Sv increase in overturning, the heat transport increases by 0.1PW. For random selections of parameter sets from the prior, the maximum overturning ranges from 0 to 30Sv, and the heat transport varies from less than 0.4 to greater than 1.6PW, so the results plotted here are not an intrinsic property of the model but are instead determined through the assimilation process. The width of our ensemble also seems comparable to the uncertainty in the observations, in contrast to the much greater range of the CMIP models.

The position of the maximum overturning varies somewhat between ensemble members, with most close to 25°N but some being positioned around 40°N where there is a small local maximum in the ensemble mean.

Estimates of the northward heat transport around 30°S range from 0.16–0.68PW (Bennett, 1978), with several recent estimates being very close to 0.3PW with a mass

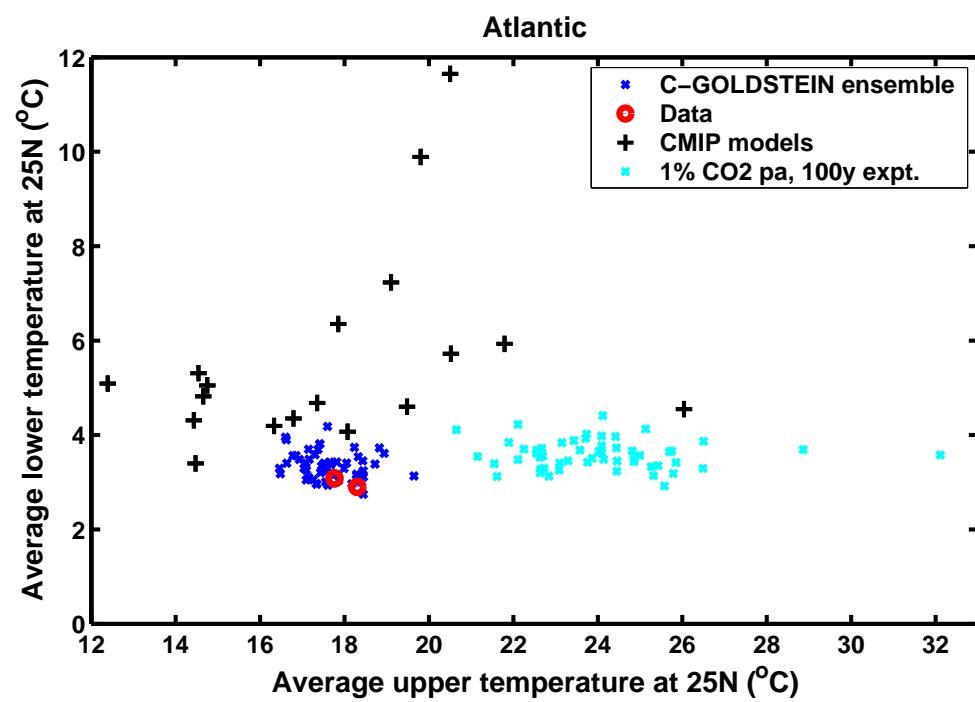
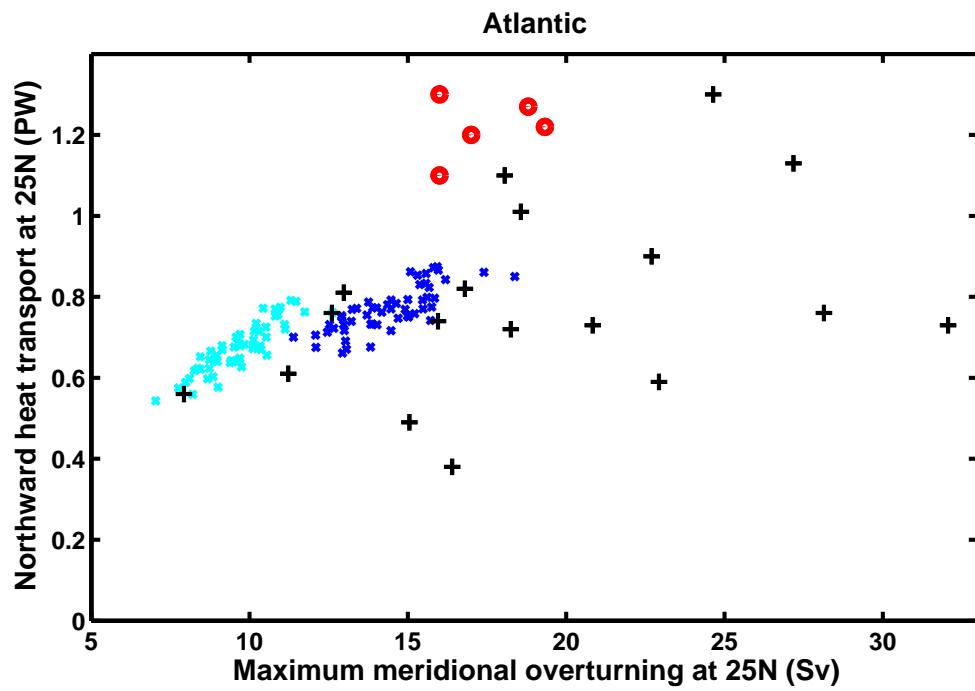


Figure 7: North Atlantic heat and mass transports, and mean temperatures of northward-flowing (upper) and southward-flowing (lower) branches of overturning circulation

transport of around 14Sv Weijer *et al.* (1999); Ganachaud & Wunsch (2000); Holfort & Siedler (2001). A few of the CMIP models give negative values here, but most are in good agreement, and their overturning strengths are also reasonable, ranging between 10 and 20Sv. We also have a very small poleward heat transport of  $0.1 \pm 0.05$ PW (Figure 6), with our overturning of 8Sv again being on the low side. This weakness in overturning may be due to the lack of a warm water path from the Pacific via the Indonesian Throughflow and Agulhas leakage, which cannot be adequately resolved at the low resolution we are using here.

### 3.4.2 Pacific Ocean

In the Pacific at  $25^\circ\text{N}$ , we have observational estimates of northward heat transport of  $0.76 \pm 0.3$ PW from Bryden *et al.* (1991), and  $0.5 \pm 0.1$ PW from Ganachaud & Wunsch (2000). The CMIP models are generally a little lower than these estimates, but still consistent with them. The overturning at  $30^\circ\text{S}$  Pacific is very uncertain in the CMIP models, ranging from 4–30Sv. Our ensemble has northward heat transport of  $0.6 \pm 0.1$ PW, and our ensemble overturning of 16Sv is in the middle of the range of CMIP results.

The observed heat transport at  $30^\circ\text{S}$  is very uncertain, with there being no consensus even on its sign. Ganachaud & Wunsch (2000) suggest 0.6PW northwards, whereas Trenberth *et al.* (2001) have a figure of 0.9PW southwards. The CMIP results have an even greater range, with mass transports from 2Sv southwards to 0.64Sv northwards. Our results here are harmed by the inability of our model to resolve the Indonesian Throughflow, by which, according to the analysis of Ganachaud & Wunsch (2000), 16Sv of water carries 1.4PW of heat from the Pacific to Indian oceans. As a result, our estimate of  $1.3 \pm 0.1$ PW southwards is rather large.

### 3.4.3 Indian Ocean

At  $30^\circ\text{S}$  in the Indian Ocean, Ganachaud & Wunsch (2000) indicate a heat transport of 1.5PW southwards, and the CMIP models are split between those with substantial transports in excess of 1PW, and those with very small transports of below 0.4PW (even in models which do resolve the Indonesian Throughflow). The heat flux across  $30^\circ\text{S}$  in our Indian Ocean is about 1PW southwards, even without the help of the Indonesian Throughflow.

### 3.4.4 Global ocean heat transport

The total global northward heat transport in the ocean has been estimated to be 2.0PW at  $18^\circ\text{N}$  (Trenberth *et al.*, 2001), and 1.8PW at  $25^\circ\text{N}$  (Ganachaud &

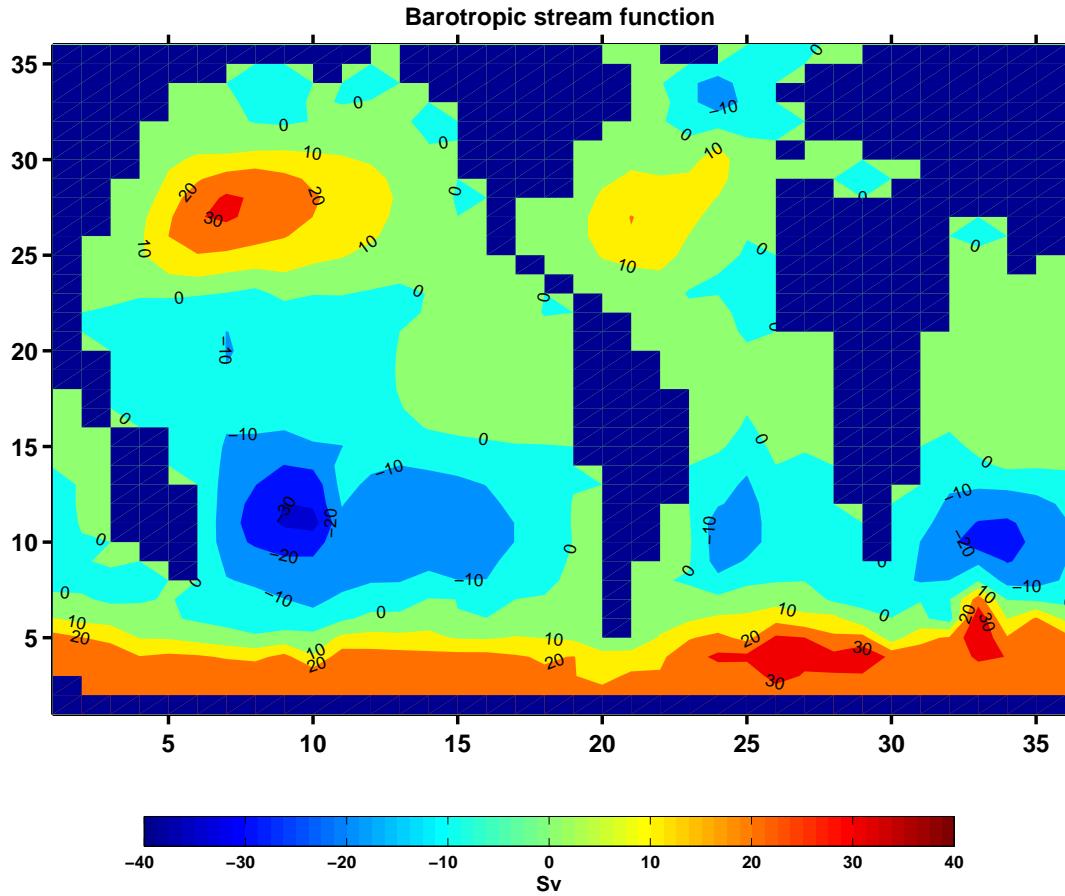


Figure 8: Barotropic circulation

Wunsch, 2000). In contrast, only 3 of the CMIP models reach 2.0PW northwards transport, and one of those includes an anomalous northward head transport in the northern Indian ocean, contrary to the sparse observational data there. The meridional heat transport of our ensemble is shown in Figure 6. The Atlantic, Pacific and total ocean heat transports are shown. Our result here of  $1.5 \pm 0.1$ PW total ocean northwards transport (at both latitudes) is good compared to most CMIP model results, although still a little low compared to the observational estimates.

The southward heat transport peaks at 1.3PW at around  $15^\circ$ S according to Trenberth *et al.* (2001). Most of the CMIP transports are substantially smaller than this, although three of the models exceed 2PW. Although the lack of Indonesian Throughflow in our model is clearly responsible for a large error in our Pacific and Indian ocean estimates at  $30^\circ$ S, the total ocean transport at this latitude (which should be less strongly affected by the exchange) is still rather too high at 2.3PW.

Barotropic flow (Figure 8) in the Atlantic and Pacific appear reasonable in our

results, similar on large scales to the results from the CCCMA model shown by Jia (2003). Of course the smaller scales are not resolved on our coarse grid. The Antarctic circumpolar current (ACC) was estimated at  $123 \pm 11$  Sv by Whitworth & Peterson (1985) and  $140 \pm 6$  Sv by Ganachaud & Wunsch (2000). The ACC is notoriously difficult to model accurately even at high resolution, with the CMIP models producing a range of results from 10—270 Sv. Our estimate is rather low, at 28 Sv, but is well within the range of the CMIP models.

### 3.5 Atmospheric state

We now consider the atmospheric state, in rather less detail since this module is very simplified and barely attempts to represent the atmospheric circulation. The zonally averaged temperature and humidity are shown in Figure 9. Although the globally integrated precipitation is reasonable, the zonally averaged precipitation (Figure 10) does not adequately represent the spatial variability in the data (in particular, it is the difference between precipitation and evaporation that affects the meridional salinity gradient), and we believe that this is largely responsible for the poor ocean salinity distribution mentioned in Section 3.3. The zonally averaged precipitation minus evaporation in our model rarely exceeds 0.2 m/year, whereas the SOC climatology (Josey *et al.*, 1998) has a peak of around 1 m/year in the equatorial region, and large troughs and peaks at  $\pm 20^\circ$  and  $\pm 45^\circ$  respectively for both hemispheres. These peaks and troughs are closely aligned with the observed zonally averaged sea surface salinity profile. The importance of this moisture transport in determining the ocean salinity was demonstrated by Weaver *et al.* (2001), when they compared model runs with moisture advection switched on and off. Our results appear comparable to their non-advection model version. Even though we have implemented an advective term here it does not appear to have sufficient effect, perhaps due to a lack of resolution and the simplified precipitation and land-surface schemes. The lack of seasonality (with the associated zonal shift of the ITCZ and its precipitation) may also be significant here.

Trenberth & Caron (2001) have examined atmospheric reanalyses from NCEP, and calculated that the northward heat transport in the atmosphere peaks at  $5.0 \pm 0.14$  PW, greatly in excess of previous estimates. Our result of  $5.1 \pm 0.2$  PW shows good agreement with their figure. The estimate from the prior was about  $6.5 \pm 1$  PW, indicating a lower dependence on the parameters and a greater degree of predetermination than was apparent for the more dynamically active ocean model. In the southern hemisphere, the observed transport is slightly higher, but our ensemble's atmospheric heat transport is slightly lower. It is not clear whether this is the cause, or the result, of the excessive heat transport in this region of the ocean. In any case, we hope for a more realistic atmospheric climatology (and therefore an improved ocean state) when a more complete 3D atmosphere is

implemented.

### 3.6 Other state variables

The only remaining component is the sea ice, which exists in each of our ensemble members in the northern hemisphere and also in the southern hemisphere of some members. Due to the equal-area grid arrangement, the northern hemisphere ice is mostly restricted to the northernmost row of cells (covering the latitudes 71–90°N) with some spread to the second row (63–71°N), and the longitudinal distribution of thickness is rather smooth. Nevertheless, the pattern and depth of ice is not unreasonable, with a maximum depth in our ensemble members of  $12 \pm 2$ m at around 90°W comparable to output from the 1/4°-resolution ocean model OCCAM (Saunders *et al.*, 1999) which has a maximum ice thickness of 7m occurring at the same longitude (Y. Aksenov, personal communication). In the southern hemisphere, the southernmost row of grid boxes is entirely occupied by land. In reality, the sea ice around the Antarctic is highly seasonal, so it seems reasonable that our ensemble sits on the boundary between being totally ice-free (31 of the 54 ensemble members), and having permanent ice cover in part of the second row of grid cells (23 of the ensemble members).

### 3.7 Summary of validation

The ensemble mean state is a reasonable one, especially in the ocean, but less realistic in the simple atmosphere. The ocean circulation is a little weak, which results in transports being generally low, although the lack of Indonesian Throughflow must also be responsible for substantial errors in the Pacific and Indian Ocean basins. It is in fact straightforward to improve the ensemble’s estimates of circulation somewhat by directly assimilating the observational estimates of its strength (not shown here). However, this procedure runs the risk of over-using data since many of these estimates are themselves modelling analyses relying to a greater or lesser extent on assimilation of temperature and salinity data. Furthermore, this would prevent us from using these observational estimates as independent validation of the model. In any case, the poor sea surface salinity distribution points to the atmospheric model being a dominant source of error. Overall, the mean state (especially in the ocean) does not seem clearly worse than the CMIP models, with each of these models also having its own particular strengths and weaknesses in different regions. Given the extremely low resolution and simplicity of our model, this is in itself an encouraging illustration of the value of the tuning method.

The ensemble width is also generally realistic, comparing well (within a factor of two or so) with the estimated uncertainty in the observational analyses such

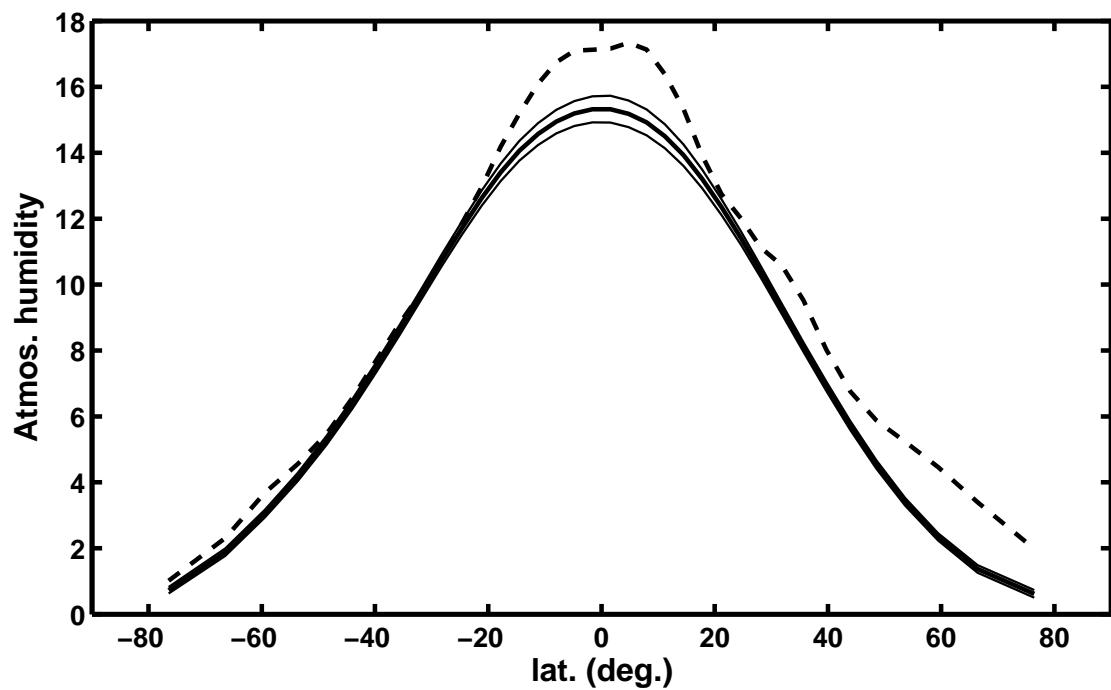
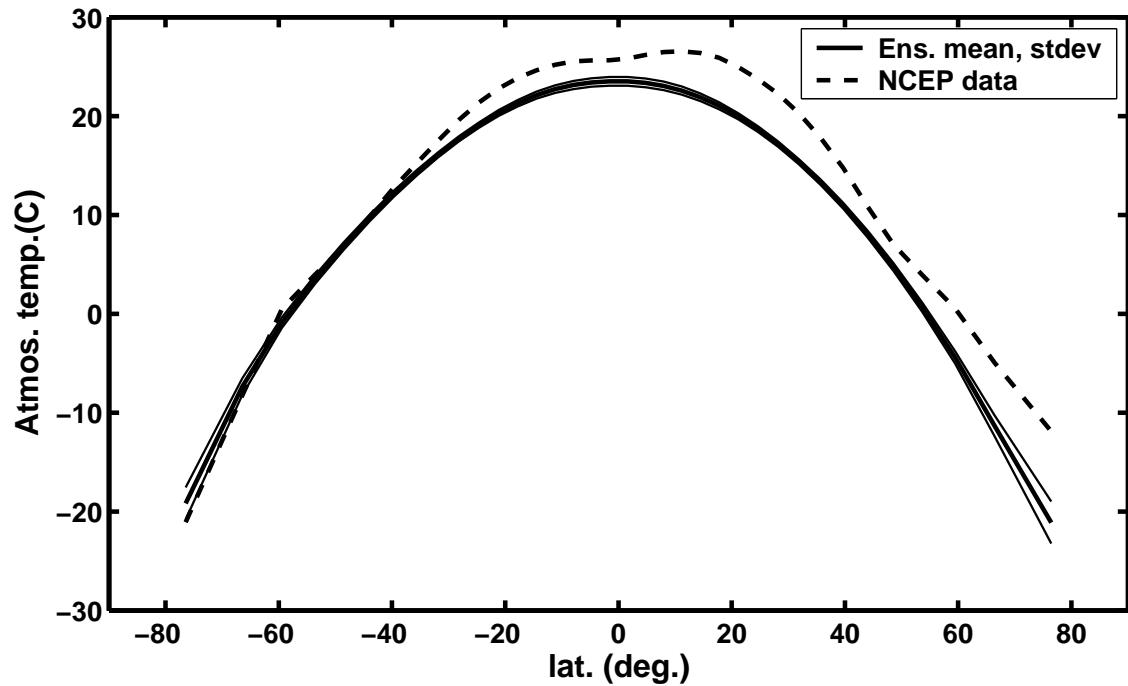


Figure 9: Zonally averaged atmospheric temperature and humidity, ensemble and data

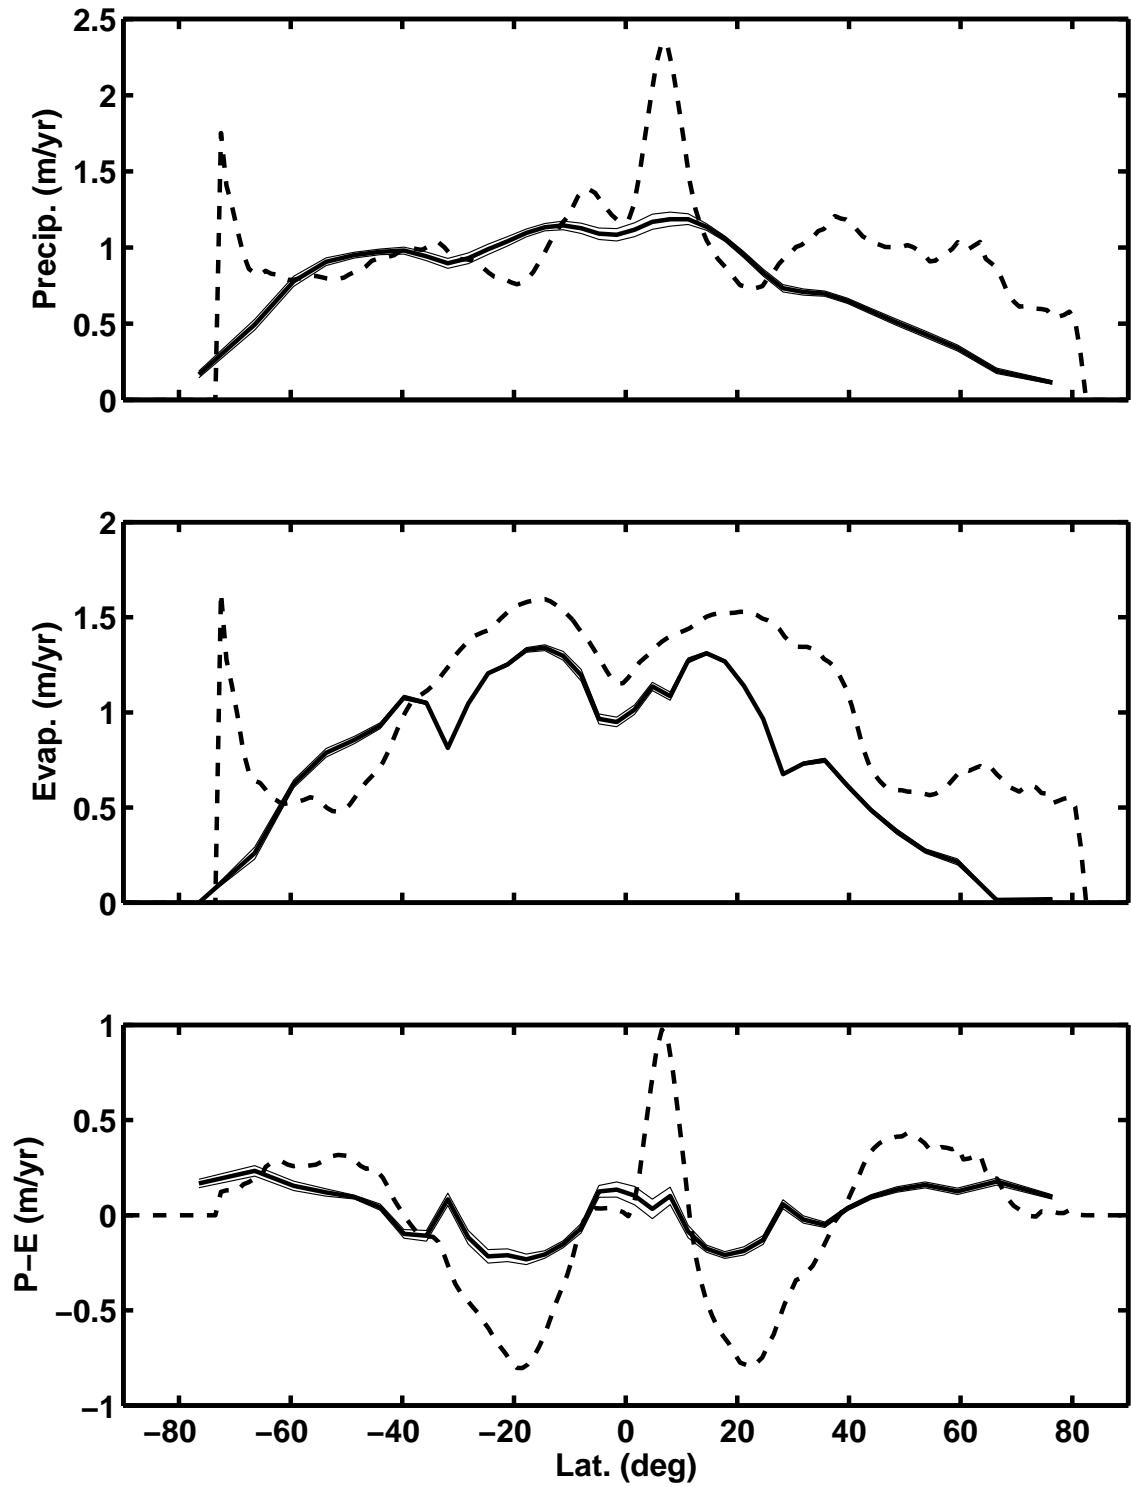


Figure 10: Zonally averaged precipitation and evaporation, ensemble mean and standard deviation, together with data from the SOC climatology

as those generated by Ganachaud & Wunsch (2000). Even when there is a bias in the mean, such as the North Atlantic overturning circulation, the ensemble still has a plausible width. This justifies our choice of scaling of the observational error estimates (which directly determines the ensemble width). Since the regional biases cannot be eliminated by tuning parameters, the model can only be improved further through more fundamental structural changes such as implementing a more realistic atmosphere. In regions where there is a large bias, the model state should not be considered as a direct estimate of the true climate state, but rather the estimated projection of the true climate state onto the space spanned by the model, and the bias should be taken into account when analysing forecasts. Our ensemble width contrasts strongly with the CMIP results, which generally span a range well in excess of any reasonable observational uncertainty, such as their estimates of maximum overturning in the North Atlantic which range from less than 8Sv to more than 32Sv. It seems improbable that an ‘ensemble of opportunity’ made up of such a wide range of models will be able to give a quantitatively useful probabilistic prediction, since the uncertainty in even the nowcast bears such little relation to the uncertainty in the true climate state.

## 4 Climate change projections

### 4.1 Forecast Scenario

Here we present some projections of climate change under a simple anthropogenic forcing scenario. In this experiment, the ensemble generated by the previous parameter tuning procedure was integrated for 70, 100 or 200 years under a 1% per annum atmospheric CO<sub>2</sub> increase (reaching 2, 2.7 or 7.3× the present day level) followed by a further 3500 years at the constant higher level. This is then followed by a slow decline in CO<sub>2</sub> at a rate of 0.05% until the value returns to the present day, and then finally a further 6000 years at the constant lower value.

In these experiments we use a fixed radiative forcing for doubled CO<sub>2</sub> of 4Wm<sup>-2</sup>, thus ignoring one source of uncertainty. The range of 3.5–4.1Wm<sup>-2</sup> of Houghton *et al.* (2001, Chapter 6) is generally taken to indicate the 95% confidence (2 $\sigma$ ) interval, so this is a rather minor factor. More importantly, this simple model does not contain any of the feedbacks related to cloud cover, which are highly uncertain but potentially large.

### 4.2 Global temperature change

The global mean surface air temperature at the start of the experiment was 12.1 ± 0.1°C, and this rose to 13.7 ± 0.1°C after 70 years of rising CO<sub>2</sub>. The mean heat

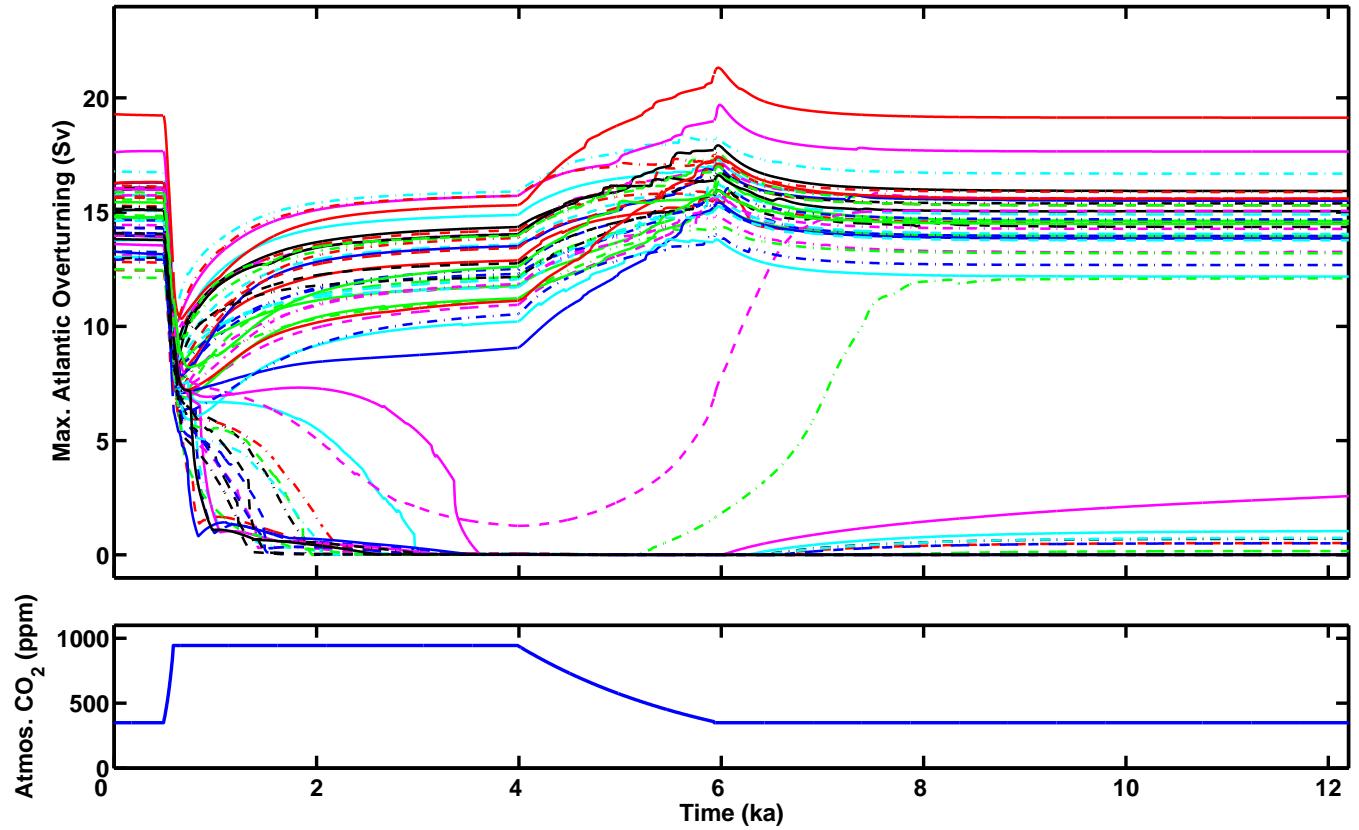


Figure 11: Evolution of the thermohaline circulation under anthropogenic forcing

rise per ensemble member was  $1.6 \pm 0.04^\circ\text{C}$ . However, after integrating for a further 5000 years to equilibrium at the  $2 \times \text{CO}_2$  level, the temperature rose further to  $15.1 \pm 0.1^\circ\text{C}$ , a final ‘climate sensitivity’ (equilibrium response to doubled  $\text{CO}_2$ ) of  $2.9 \pm 0.1^\circ\text{C}$ . The integrations with the longer warming periods showed that the equilibrium temperature response is roughly linear in radiative forcing. The uncertainties on all these figures are markedly lower than those estimated by other researchers, probably due to the lack of competing but highly uncertain feedbacks in the atmospheric model (eg cloud effects). However, it should be noted that our initial ensemble has a reasonable range of uncertainty on most of the state diagnostics examined earlier, whereas the ensemble of coupled models in the CMIP experiments have a huge range of uncertainty which is generally much greater than that of the observations. Therefore, it seems possible that the range of predictions generated as an ‘ensemble of opportunity’ from these models is unrealistically wide. In any case, our results suggest that the use of a ‘climate sensitivity’ estimate needs to be carefully qualified in terms of the time allowed to approach the new equilibrium. Higher resolution models cannot realistically be integrated for such long periods as are necessary for a true equilibrium (which in any case may have limited utility), but the state reached during or immediately following a transient rise in  $\text{CO}_2$  will be far from the equilibrium.

### 4.3 THC collapse and hysteresis

Recently, the strength and stability of the North Atlantic thermohaline circulation (THC) has been a subject of much interest. It is widely believed that the THC has two or more quasi-stable states, and has repeatedly switched abruptly between the current ‘on’ state, and substantially a weaker, or even ‘off’ state over the recent paleoclimate record (Broecker, 1997). Since the THC is responsible for transporting a large amount of heat to the North Atlantic and Western Europe, such transitions are accompanied by large regional changes in climate. The possibility of an abrupt transition has therefore been the subject of intense study in recent years (eg Broecker, 1997; Rahmstorf & Ganopolski, 1999). Many studies have confirmed the sensitivity of the THC in ocean models to externally imposed freshwater fluxes, and more recently, coupled models have been used to investigate the response of the atmosphere-ocean system to anthropogenically enhanced  $\text{CO}_2$  (Wiebe & Weaver, 1999; Stouffer & Manabe, 2003). The results generally indicate that the THC can be switched off by a small additional freshwater flux of around 0.06–0.15Sv. This figure is comparable to some estimates of likely changes in the hydrological cycle, although there is substantial disagreement between different models regarding the overall effect, with some indicating a weakening or even complete shutdown (Cubasch *et al.*, 2001), and others a strengthening due to increased freshwater flux further south (Latif *et al.*, 2000). Moreover, in some models where the THC initially shuts down, it is re-established in the longer term

(under steady but elevated CO<sub>2</sub> levels) to a value which may even exceed the original state (Wiebe & Weaver, 1999; Stouffer & Manabe, 2003). The wide range of results generated by different models indicates a substantial uncertainty over the response of the THC to anthropogenic perturbations, over both the short and longer term.

The time series of overturning in the North Atlantic (for the experiment using 100 years of CO<sub>2</sub> increase), for each ensemble member, is shown in Figure 11. The initial response to the rise in CO<sub>2</sub> is a rapid drop in the strength of the THC. The cyan crosses in Figure 7 show the snapshot of overturning and temperatures in the North Atlantic at the end of the 100 year period. Over this initial interval, the surface warming almost completely compensates for the reduction in overturning strength, with the total heat transport only dropping marginally. During the stabilisation phase of the experiment, the overturning largely recovers in many of the ensemble members, but continues to fall and ultimately collapses in a substantial proportion of them (19 out of 54, or 37%). This results in a substantial difference in the regional climate, with the temperature in NW Europe being around 3-5°C colder in those ensemble members where the THC collapses, when compared to those in which it recovers. In fact, the regional temperature for collapsed members is generally colder than the present day climate, but only by 1-2°C. The time taken for the collapse ranges from around 250 years to 3000 years. Even in those members where the overturning recovers somewhat, it remains at a generally lower level than initially, in contrast to the results of Wiebe & Weaver (1999) and Stouffer & Manabe (2003). Our simple atmosphere is incapable of substantially varying the moisture transport from the Atlantic to the Pacific basins, but the model's dynamically diagnosed moisture transport does drop by around 0.01–0.02Sv which adds to the surface warming effect in helping to destabilise the THC. Using another efficient model with simple atmosphere, Rahmstorf & Ganopolski (1999) added a further external perturbation to the freshwater forcing in proportion to the change in the model's northern hemisphere temperature. However, the constant of proportionality for this perturbation is unknown *a priori* and cannot be estimated even in principle from a single steady state climatological tuning, since their parameterisation is defined in terms of the deviation from the climatological mean. Therefore, we have not attempted to include this factor in our study, although it could have a significant influence on THC stability.

During and after the atmosphere's return to present-day CO<sub>2</sub> levels, the THC increases back towards the initial level in those models where it did not collapse entirely. Of the ensemble members where the THC switched off, only two switch back on again, illustrating the hysteresis which has been widely investigated by others. When a similar global warming experiment was performed with 70 years of 1% pa CO<sub>2</sub> increase, only 6 of the models switched to the collapsed THC state (11% of the ensemble). With 200 years of increase, 52 of them collapsed

(96%). These results, which are the first using a truly objective multivariate analysis system to simultaneously sample the uncertainty of the climate system due to many different parameterisations, suggest that the current state of the THC could be vulnerable to anthropogenic perturbation, and if atmospheric CO<sub>2</sub> levels increase rapidly, we could see drastic changes in the thermohaline circulation in the next couple of centuries. However it should be noted that the subset of ensemble members that shut down their overturning are predominantly those with the lower initial values, and since the ensemble as a whole has a low bias, the real risk may be somewhat overstated by these results. Missing processes in the model (in particular, the inadequacy of the moisture transport and the use of a fixed wind field) also limits the confidence that can be placed in these quantitative results. However, application of this methodology to more realistic models should help to quantify the risk more accurately.

## 5 Conclusions

We have applied the method of Annan *et al.* (2004) to perform probabilistic multivariate parameter estimation by assimilating observational data for the ocean and atmosphere into a new highly efficient coupled global atmosphere-ocean model. The method generates an ensemble whose members sample the uncertainty of the current climate state. The ensemble mean appears to have a very reasonable steady state climatology (especially in the ocean, which is the most sophisticated component of the model), within the constraints of the model's limited physics and resolution. In fact, it appears comparable in realism to the those of the much more complex and expensive coupled models used in the CMIP project (Jia, 2003). Furthermore, the ensemble spread is comparable to the uncertainty estimated by other ocean state analyses (Ganachaud & Wunsch, 2000, eg). In principle, such a tuned ensemble should be useful for making objective predictions of future climate change under anthropogenic forcing. However, it is clear that the simplicity of the atmospheric component of this model limits its value for this purpose.

When tuned to present day climatology, the forecast under a scenario of 70 years of steadily rising atmospheric CO<sub>2</sub> (1% per annum cumulative growth) is for a rise in surface air temperature of 1.6°C in 70 years, increasing to 2.9°C at equilibrium. These results are consistent with other research. However, this model is rather simplistic (particularly with respect to the atmosphere) and therefore the importance of these results may be considered to be more in terms of what the method promises for more sophisticated models (including, but not limited to, future versions of the GENIE model) than in terms of their accuracy for the real Earth climate system.

The North Atlantic overturning is of particular interest. It is widely (but not unan-

imously) believed that atmospheric CO<sub>2</sub> increase is likely to result in a reduction in overturning. Recent model results have suggested that, over the longer term, overturning will recover and indeed converge to a slightly higher strength than the original value, although there are some differences in detail between the different model results (Wiebe & Weaver, 1999; Stouffer & Manabe, 2003). A further unknown is the level required for a permanent (or at least long-term) shutdown of the overturning, which is considered a likely result at some high level of CO<sub>2</sub>. Our results indicate a polarisation of the ensemble into two classes of behaviour: those where a recovery takes place, and those where a complete shutdown occurs for our 2.7× CO<sub>2</sub> increase. When this experiment was repeated for a 200 year CO<sub>2</sub> increase and stabilisation (7.3×), only 2 of the 54 ensemble members remained in the ‘on’ state. These results suggest a significant (if somewhat unquantified) risk of substantial and effectively irreversible changes in regional climate in response to anthropogenic perturbation.

## 6 Acknowledgements

NRE is currently supported by the Swiss NCCR-Climate programme. RM is supported by the Natural Environment Research Council (NERC) Core Strategic Programme “Ocean Variability and Climate”. Supercomputer facilities and support were provided by JAMSTEC. This research was partly supported by the GENIE project (<http://www.genie.ac.uk/>), which is funded by the NERC (NER/T/S/2002/00217) through the e-Science programme.

## References

Andronova, N. G. & M. E. Schlesinger, 2001: Objective estimation of the probability density function for climate sensitivity, *Journal of Geophysical Research*, **108**(D8), 22,605–22,611.

Annan, J. D. & J. C. Hargreaves, 2004: Efficient parameter estimation for a highly chaotic system, *Tellus*, in press.

Annan, J. D., J. C. Hargreaves, N. R. Edwards & R. Marsh, 2004: Parameter estimation in an intermediate complexity Earth System Model using an ensemble Kalman filter, *Ocean Modelling*, in press.

Bennett, A. F., 1978: Poleward heat fluxes in Southern Hemisphere Oceans, *Journal of Physical Oceanography*, **8**, 785–798.

Böning, C. W., F. O. Bryan, W. R. Holland & R. Döscher, 1996: Deep water formation and meridional overturning in a high-resolution model of the North Atlantic, *Journal of Physical Oceanography*, **26**, 1142–1164.

Broecker, W. S., 1997: Thermohaline circulation, the Achilles heel of our climate system: will man-made CO<sub>2</sub> upset the current balance?, *Science*, **278**, 1582–1588.

Bryden, H. L., D. H. Roemich & J. A. Church, 1991: Ocean heat transport across 24°N in the Pacific, *Deep-Sea Research*, **48**, 297–324.

Collins, M. & M. R. Allen, 2002: Assessing the relative roles of initial and boundary conditions in interannual to decadal climate predictability, *Journal of Climate*, **15**(21), 3104–3109.

Cubasch, U., A. Meehl, G. J. Boer, R. J. Stouffer, M. Dix, A. Noda, C. A. Semenov, S. Raper & K. S. Rap, 2001: Projections of future climate change, in J. T. Houghton, Y. Ding, D. J. Griggs, M. Nouger, P. J. v. d. Linden, X. Dai, K. Maskell & C. A. Johnson, editors, *Climate Change 2001: The Scientific Basis. Contribution of Working Group 1 to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, pages 525–582, Cambridge University Press.

Edwards, N. R. & R. Marsh, 2003: Uncertainties due to transport-parameter sensitivity in an efficient 3-D ocean-climate model, *Climate Dynamics*, submitted.

Edwards, N. R. & J. G. Shepherd, 2002: Bifurcations of the thermohaline circulation in a simplified three-dimensional model of the world ocean and the effects of interbasin connectivity, *Climate Dynamics*, **19**, 31–42.

Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *Journal of Geophysical Research*, **99**(C5), 10143–10162.

Forest, C. E., M. R. Allen, P. H. Stone & A. P. Sokolov, 2000: Constraining uncertainties in climate models using climate change detection techniques, *Geophysical Research Letters*, **27**, 569–572.

Fukumori, I., 2001: *Data assimilation by models*, chapter 5, pages 237–265, Academic Press.

Ganachaud, A. & C. Wunsch, 2000: Improved estimates of global ocean circulation, heat transport and mixing from hydrographic data, *Nature*, **408**, 453–457.

Gregory, J. M., R. J. Stouffer, S. C. B. Raper, P. A. Stott & N. A. Rayner, 2002: An observationally based estimate of the climate sensitivity, *Journal of Climate*, **15**(22), 3117–3121.

Hansen, J., M. Sato, L. Nazarenko, R. Ruedy, A. Lacis, D. Koch, I. Tegen, T. Hall, D. Shindell, P. Stone, T. Novakov, L. Thomason, R. Wang, Y. Wang, D. Jacob, S. Hollandsworth-Frith, L. Bishop, J. Logan, A. Thompson, R. Stolarski, J. Lean, R. Willson, S. Levitus, J. Antonov, N. Rayner, D. Parker & J. Christy, 2002: Climate forcings in Goddard Institute for Space Studies SI2000 simulations, *Journal of Geophysical Research*, **107**(D18), 4347–4384.

Hargreaves, J. C. & J. D. Annan, 2002: Assimilation of paleo-data in a simple Earth system model, *Climate Dynamics*, **19**(5–6), 371–381.

Hibler, W. D., 1979: A dynamic thermodynamic sea ice model, *Journal of Physical Oceanography*, **9**, 815–846.

Holfort, J. & G. Siedler, 2001: The meridional oceanic transports of heat and nutrients in the South Atlantic, *Journal of Physical Oceanography*, **31**, 5–29.

Houghton, J. T., Y. Ding, D. J. Griggs, M. Noguer, P. J. V. D. Linden, X. Dai, K. Maskell & C. A. Johnson, 2001: *Climate Change 2001: Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press.

Jia, Y., 2003: Ocean heat transport and its relationship to ocean circulation in CMIP, *Climate Dynamics*, **20**, 153–174.

Josey, S. A., E. C. Kent & P. K. Taylor, 1998: The Southampton Oceanography Centre (SOC) ocean — atmosphere heat, momentum and freshwater flux atlas, Technical Report 6, Southampton Oceanography Centre.

Keppenne, C. L., 2000: Data assimilation into a primitive-equation model with a parallel ensemble Kalman filter, *Monthly Weather Review*, **128**, 1971–1981.

Knutti, R., T. F. Stocker, F. Joos & G.-K. Plattner, 2002: Constraints on radiative forcing and future climate change from observations and climate model ensembles, *Nature*, **416**, 719–723.

Lambert, S. J. & G. J. Boer, 2001: CMIP1 evaluation and intercomparison of coupled climate models, *Climate Dynamics*, **17**, 83–106.

Latif, M., E. Roeckner, U. Mikolajewicz & R. Voss, 2000: Tropical stabilisation of the thermohaline circulation in a greenhouse warming simulation, *Journal of Climate*, **13**, 1809–1813.

Levitus, S., 1998: *World Ocean Atlas*, <http://www.cdc.noaa.gov/>.

Levitus, S., J. I. Antonov, T. P. Boyer & C. Stephens, 2000: Warming of the world ocean, *Science*, **287**, 2225–2229.

Oort, A. H., 1983: Global atmospheric circulation statistics, 1958–1973, NOAA Professional Paper 14.

Press, W. H., S. A. Teukolsky, W. T. Vetterling & B. P. Flannery, 1994: *Numerical recipes in Fortran: the art of scientific computing*, Cambridge University Press.

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

Saunders, P. M., A. C. Coward & B. A. de Cuevas, 1999: Circulation of the Pacific Ocean seen in a global ocean model (OCCAM), *Journal of Geophysical Research*, **C8**, 18281–18299.

Semtner, A. J., 1976: A model for the thermodynamic growth of sea ice in numerical investigations of climate, *Journal of Physical Oceanography*, **6**, 379–389.

Stouffer, R. J. & S. Manabe, 2003: Equilibrium response of thermohaline circulation to large changes in atmospheric CO<sub>2</sub> concentration, *Climate Dynamics*, **20**, 759–773.

Thacker, W. C., 2003: Data-model-error compatibility, *Ocean Modelling*, **5**, 233–247.

Trenberth, K. E. & J. M. Caron, 2001: Estimates of meridional atmosphere and ocean heat transports, *Journal of Climate*, **14**, 3433–3443.

Trenberth, K. E., J. M. Caron & D. P. Stepaniak, 2001: The atmospheric energy budget and implications for surface fluxes and ocean heat transports, *Climate Dynamics*, **17**, 259–276.

Weaver, A. J., M. Eby, E. C. Wiebe, C. M. Bitz, P. B. Duffy, T. L. Ewen, A. F. Fanning, M. M. Holland, A. MacFadyen, H. D. Matthews, K. J. Meissner, O. Saenko, A. Schmittner, H. X. Wang & M. Yoshimori, 2001: The UVic Earth System Climate Model: Model description, climatology, and applications to past, present and future climates, *Atmosphere–Ocean*, **39**(4), 361–428.

Weijer, W., W. P. M. de Ruijter, H. A. Dijkstra & P. J. van Leeuwen, 1999: Impact of interbasin exchange on the Atlantic overturning circulation, *Journal of Physical Oceanography*, **29**, 2266–2284.

Whitworth, T. & R. G. Peterson, 1985: Volume transport of the Antarctic Circumpolar Current from bottom pressure measurements, *Journal of Physical Oceanography*, **15**, 810–816.

Wiebe, E. C. & A. J. Weaver, 1999: On the sensitivity of global warming experiments to the parameterization of sub-grid scale mixing, *Climate Dynamics*, pages 875–893.