

Decadal predictability of the Atlantic Ocean in a coupled GCM: forecast skill and optimal perturbations using Linear Inverse Modelling

ED HAWKINS* AND ROWAN SUTTON
WALKER INSTITUTE, DEPARTMENT OF METEOROLOGY,
UNIVERSITY OF READING

ACCEPTED BY JOURNAL OF CLIMATE
doi:10.1175/2009JCLI2720.1

February 10, 2009

ABSTRACT

The decadal predictability of three-dimensional Atlantic Ocean anomalies is examined in a coupled global climate model (HadCM3) using a Linear Inverse Modelling (LIM) approach. It is found that the evolution of temperature and salinity in the Atlantic, and the strength of the meridional overturning circulation (MOC), can be effectively described by a linear dynamical system forced by white noise. The forecasts produced using this linear model are more skillful than other reference forecasts for several decades. Furthermore, significant non-normal amplification is found under several different norms. The regions from which this growth occurs are found to be fairly shallow and located in the far North Atlantic. Initially, anomalies in the Nordic Seas impact the MOC, and the anomalies then grow to fill the entire Atlantic basin, especially at depth, over one to three decades. It is found that the structure of the optimal initial condition for amplification is sensitive to the norm employed, but the initial growth seems to be dominated by MOC-related basin scale changes, irrespective of the choice of norm. The consistent identification of the far North Atlantic as the most sensitive region for small perturbations suggests that additional observations in this region would be optimal for constraining decadal climate predictions.

* *Corresponding author address:* Ed Hawkins, Department of Meteorology, University of Reading, Reading, RG6 6BB. UK. E-mail: e.hawkins@reading.ac.uk

1. Introduction

Climate variability on decadal timescales is strongly influenced by the oceans. The Atlantic Ocean is particularly important because of its central role in the overturning circulation. Furthermore, there is evidence from models that variations in the Atlantic overturning circulation, and associated impacts on climate, are potentially predictable on decadal timescales (e.g. Griffies and Bryan 1997; Collins and Sinha 2003; Pohlmann et al. 2004; Sutton and Hodson 2005; Collins et al. 2006). This evidence suggests that predictions for the climate of the next few decades should be properly initialised using information about the present ocean state, rather than relying solely on the simulated response to changing radiative forcing (Solomon et al. 2007).

There has recently been significant progress in the development of properly initialised decadal prediction systems (Smith et al. 2007; Keenlyside et al. 2008; Meehl et al. 2008). Nevertheless, such systems are still at an early stage of development, and there are many challenges ahead. One important need is to improve understanding of the processes that govern error growth. As in numerical weather prediction, understanding the growth of perturbations is essential in order to design ensembles for sampling uncertainty efficiently (e.g. Molteni et al. 1996), and also to improve observing systems in a targeted, cost-effective way. This latter point is especially important for the oceans because of the high cost of sub-surface observations.

The purpose of our study is to investigate the growth, on decadal timescales, of perturbations in the Atlantic Ocean. More specifically, we aim to estimate optimal perturbations, these being the perturbations that grow most rapidly, in a linear sense, as quantified by an appropriate metric, over a defined time interval (e.g. Farrell 1988). These perturbations have been examined in simple models using a variety of techniques (e.g. Lohmann and Schneider 1999; Zanna and Tziperman 2005; Sevellec et al. 2008). The methodology we employ is the Linear Inverse Modelling (LIM) approach of Penland and Sardeshmukh (1995), which we apply to output from a coupled General Circulation Model (GCM). The LIM approach is attractive in that it provides a way to estimate optimal ‘non-normal’ perturbations with vastly less computational expense than alternative methods (e.g. empirical singular vectors, Kleeman et al. 2003). LIM has been exploited to study decadal variability and predictability in the Pacific using observations (e.g. Newman 2007; Alexander et al. 2008), and amplification of the Atlantic thermohaline circulation (THC) in simple box models (e.g. Tziperman and Ioannou 2002). Recently, Tziperman et al. (2008) used LIM to examine decadal predictability in the Atlantic Ocean in a coupled GCM (GFDL CM2.1) and found significant rapid, non-normal amplification from anomalies near the surface in the Labrador Sea, which grew into anomalies in a larger region of the North Atlantic. In their model, this growth limits the predictability of the ocean temperature and salinity fields, as well as the THC, to about 8 years.

However, it is important to determine if this predictability limit is found in other GCMs (and in the real ocean), and there is evidence that different GCMs behave in diverse ways.

Griffies and Bryan (1997) analysed a previous version of the GFDL model and found that the predictability of the deep ocean lasted 10-20 years. Collins and Sinha (2003) found that the THC in HadCM3 was potentially predictable for up to 50 years. This analysis was extended to a multi-model study by Collins et al. (2006) who found that several GCMs had significant predictability for 20 years. There was also evidence that predictability varies depending on the exact initial conditions, e.g. whether the overturning is relatively weak or strong.

In this paper we adopt the LIM approach to explore Atlantic Ocean predictability in HadCM3. This analysis, although similar to Tziperman et al. (2008), was developed independently and has a different emphasis. The forecast skill, and the robustness of the identified optimal perturbations to the LIM methodology, under several different norms, are explored in detail. The suggested physical mechanisms for amplification identified by the linear model are shown to occur in the full GCM. These robustness checks are vital for the successful application of these results to operational decadal climate predictions, and to the design of future ocean monitoring systems.

This paper is structured as follows. In Section 2 we briefly describe the HadCM3 model and the data used. Section 3 summarises the use of LIM, and the justification for this approach, with more details in Appendices A and B. Section 4 demonstrates the ability of the constructed linear model to produce skillful forecasts for several decades. The estimation of the non-normal amplification and optimal initial conditions is presented in Section 5, and we conclude and discuss the implications for predictability in Section 6.

2. Model description and data used

In the subsequent analysis we have used data from an extended (1600 years) pre-industrial control run of the Hadley Centre climate model, HadCM3. The model details are given in Gordon et al. (2000) and references therein, and here we give a brief summary. HadCM3 is a global coupled ocean-atmosphere model with an atmospheric resolution of $2.5^\circ \times 3.75^\circ$ and 19 vertical levels. The ocean component has a resolution of $1.25^\circ \times 1.25^\circ$ with 20 vertical levels. The model does not require flux adjustment to maintain a stable climate. The mean state of the ocean model matches observed values to within 1 K and 1 psu in most regions (Gordon et al. 2000; Pardaens et al. 2003). We do not analyse the first 500 years of data to minimise the influence of spin-up effects in the ocean.

a. EOF basis

Three-dimensional (3D) bivariate correlation EOFs of the annual mean data for temperature and salinity, in the Atlantic, for 1100 years, were estimated by Hawkins and Sutton (2007) (hereafter HS07), who describe the details of their construction. In summary, the domain used is $20^\circ\text{S} - 90^\circ\text{N}$ and $100^\circ\text{W} - 20^\circ\text{E}$, using twelve depth levels from the surface to 1800m. The state vector contains both salinity and temperature anomalies, which are normalised to have unit variance at each grid point, and are also weighted by their con-

tribution to local density. We also ensure that each volume of water is treated equally by weighting by the volume of each grid box. HS07 showed that the two leading modes have significant multi-decadal and centennial variability and are well correlated with an index of the meridional overturning circulation (MOC) on decadal time scales. These two modes also show good potential predictability which is explored further here.

In this study we represent the state of the Atlantic Ocean using the leading 30 EOFs, which have been scaled so that the standard deviation of each principal component (PC) is proportional to the square root of the fraction of variance of the full three-dimensional fields explained. The eigenspectrum is shown in Fig. 1 as the solid black line. The dark grey line represents the cumulative variance explained and shows that these 30 EOFs together account for 71% of the total variance of this large (105150×1100 elements) system. The leading four PCs are shown in Fig. 2.

b. Overturning index

The strength of the meridional overturning circulation (MOC) in this model has been found to vary on a wide range of timescales¹, including an inter-decadal mode (~ 25 years, Dong and Sutton 2005), a multi-decadal mode (~ 70 years, Vellinga and Wu 2004) and a centennial mode (~ 150 years, HS07). In this paper we will examine whether small stochastic perturbations can amplify and excite MOC variability.

The overturning index used in this study is defined as the anomaly from the time mean meridional streamfunction averaged over the latitude band $27.5^\circ\text{N} - 32.5^\circ\text{N}$, at a depth of 1000m. This is the same as used by HS07, but here it has been decadal smoothed. The mean overturning strength for this smoothed index is 16.5 Sv with a standard deviation of 0.59 Sv, and is hereafter denoted as MOI_{GCM} . The light grey line in Fig. 1 shows that the leading 30 PCs account for 81% of the variance of this index, at zero lag.

3. Linear inverse modelling

We now briefly describe the linear inverse modelling (LIM) approach, following Penland and Sardeshmukh (1995) (hereafter PS95). The more technical details are given in Appendix A.

a. The linear model

The evolution of variables, \mathbf{y} , in a GCM can be represented as,

$$\frac{d\mathbf{y}}{dt} = F(\mathbf{y}), \quad (1)$$

where F is a complex non-linear operator. We assume here that the dynamics are effectively linear, i.e. that the non-linear dynamics have a much shorter timescale than the linear

¹The power spectrum of the MOC variability is shown later in Fig. 14.

dynamics (Penland 1996). The temporal evolution of the GCM can then be approximated as a stochastically forced linear dynamical system,

$$\frac{d\mathbf{x}}{dt} = \mathbf{B}\mathbf{x} + \xi, \quad (2)$$

where \mathbf{x} is a state vector, ξ is a forcing term, and \mathbf{B} is a matrix defining the temporal evolution of the state vector. We discuss later how well this linear model performs for the ocean state in HadCM3.

If $\mathbf{B}^T\mathbf{B} \neq \mathbf{B}\mathbf{B}^T$ then the system is described as non-normal, which results in eigenvectors of \mathbf{B} that are non-orthogonal. This is the case in most systems based on fluid dynamical equations such as that being considered here (e.g. Farrell 1988).

To reduce the dimensionality of the system we consider just the evolution in the subspace of the leading modes of variability in the data (in this case, the EOFs²). The matrix \mathbf{B} which best models the PCs of the GCM data can be estimated through the data covariance matrices as described in Appendix A. Properties of the noise, ξ , can be similarly estimated.

The eigenvectors of \mathbf{B} are sometimes known as the Empirical Normal Modes (ENMs, PS95), and their properties for our dataset are discussed further in Appendix A. These modes are not usually promising for examining growth as they tend to be overdamped, oscillating modes. For our dataset, one exponentially decaying mode shows a small transient amplification when initialized in certain phases, but larger growth is possible through the interference (or sum) of different modes (see Section 5).

The linear model described makes several assumptions about the linear nature of the system and the properties of the stochastic forcing. Previous authors have described several tests to justify these assumptions, and the application of these tests to our dataset is discussed in Appendix B. We restrict our analysis to the leading 30 EOFs as these tests are not well passed when more EOFs are considered.

b. Making forecasts

The matrix \mathbf{B} can be used to make forecasts of \mathbf{x} for any lead time, τ ,

$$\hat{\mathbf{x}}(t + \tau) = \mathbf{P}_\tau\mathbf{x}(t) \quad (3)$$

where $\hat{\mathbf{x}}$ denotes the predicted value of \mathbf{x} , and our general linear propagator,

$$\mathbf{P}_\tau = \exp(\tau\mathbf{B}). \quad (4)$$

For comparison, we also model the evolution of each PC (i.e. each element of \mathbf{x}) as an autoregressive (AR1) process, which uses a diagonal matrix for \mathbf{P}_τ . Results using both the full and diagonal propagators are shown later.

²Note that although Farrell and Ioannou (2001) argue that EOFs may not be the optimal choice for investigating non-normal growth, we find that they describe the system adequately for our purposes.

As well as forecasting the PC state vector, \mathbf{x} , it is also possible to forecast any index of variability. Here we focus on the overturning strength and define a vector, \mathbf{m} , which optimally combines the PCs to create a reconstructed overturning index,

$$\text{MOI}_{\text{recon}}(t) = \mathbf{m}^T \mathbf{x}(t), \quad (5)$$

by minimising the variance of the residuals between MOI_{GCM} and $\text{MOI}_{\text{recon}}$ (e.g. Tziperman et al. 2008). The reconstructed MOI index can then be forecast in a similar way to Eq. (3),

$$\widehat{\text{MOI}}_{\text{recon}}(t + \tau) = \mathbf{m}^T \widehat{\mathbf{x}}(t + \tau) = \mathbf{m}^T \mathbf{P}_\tau \mathbf{x}(t), \quad (6)$$

and compared with the actual $\text{MOI}_{\text{recon}}$ (and MOI_{GCM}) to give an estimate of the skill of the forecast.

4. Forecasts and predictability

Having demonstrated that our system is behaving appropriately (Appendix B) we can proceed to use the linear modelling approach described above with more confidence. We first consider the skill of the forecasts produced with Eqs. (3) and (6).

a. Forecasting the leading PCs

Fig. 2 shows the four leading PCs (grey lines) and forecasts for lead times of up to 50 years (black lines) using our linear model. The forecasts seem to show considerable skill, even with multi-decadal lead times. We now quantitatively assess this skill by comparing the linear model forecast errors to reference forecasts such as climatology ($\mathbf{P}_\tau = 0$) and persistence ($\mathbf{P}_\tau = 1$). We also compare with the theoretical error due to the presence of the unpredictable white noise forcing in Eq. (2) - this is termed a ‘perfect’ linear model. Penland (1989) showed that the prediction error covariance matrix, assuming the system was purely linear, would be

$$\sigma(\tau) = \mathbf{C}(0) - \mathbf{P}_\tau \mathbf{C}(0) \mathbf{P}_\tau^T \quad (7)$$

where $\mathbf{C}(0)$ is the covariance matrix of \mathbf{x} , defined in Eqn. A4. As we are fitting a linear model to a system which is probably slightly non-linear the actual errors using the propagator will be different, and normally larger. So, the differences between the actual and theoretical errors is a measure of the non-linearity of the system and is a test of our linear assumptions.

Fig. 3 shows the rms errors in the forecasts for the total vector \mathbf{x} (top left panel) and the leading five PCs for a perfect linear model (Eq. (7), dashed black line), the full linear model (solid black line), the restricted diagonal model (solid grey line, which is equivalent to a damped persistence (AR1) forecast) and a pure persistence forecast (dashed grey line), relative to a climatological forecast. The full linear model comfortably beats climatology and the diagonal model for up to 100 years in all examples shown. The full linear model is also close to the theoretical limit, demonstrating that our system is behaving roughly linearly. The largest non-linearity appears in the leading three PCs.

The rms errors shown in Fig. 3 are potentially optimistic because the errors and propagator are estimated using the same data. In Fig. 4 we test how important this is by showing the rms errors from just the last 100 years calculated using the full propagator (dashed black line) and a propagator calculated using just the first 1000 years, i.e. independent data (solid grey line). The solid black lines are repeated from Fig. 3. Although the rms errors tend to increase when using independent data, they are not dramatically worse, particularly for the first few decades. All further results are shown using the full propagator. The large differences between the solid and dashed black lines demonstrate that the rms errors can be state dependent.

b. Forecasting the overturning

We now consider forecasts of the overturning indices, which rely on the higher PCs not considered in detail above, although it should be noted that not all of the MOI variability is captured by the linear model (see Appendix B).

1) THE RECONSTRUCTED MOI

The reconstructed MOI index, defined in Eq. (5), has a correlation of 0.90 with MOI_{GCM} .³ The time series of both indices are shown in Fig. 5, which also shows forecasts of the MOI for lead times of up to 50 years (black lines) for our full linear model. Again, the forecasts show some predictive ability, especially for the low-frequency variability, even at long lead times.

In a similar way to Eq. (7), Tziperman et al. (2008) showed that the theoretical error covariance for $\text{MOI}_{\text{recon}}$, assuming a perfectly linear system, is,

$$\sigma_{\text{MOI}}(\tau) = \langle \text{MOI}_{\text{recon}}(t)^2 \rangle - \mathbf{m}^T \mathbf{P}_\tau \mathbf{C}(0) \mathbf{P}_\tau^T \mathbf{m}. \quad (8)$$

The bottom left panel of Fig. 3 shows the rms errors for the forecast of $\text{MOI}_{\text{recon}}$ for the full linear model (solid black line), the perfect linear system (dashed black line), the diagonal model (solid grey line) and the persistence forecast (dashed grey line). It can be seen that the full linear model is again close to the theoretical limit and outperforms all the reference forecasts for several decades.

2) THE GCM OVERTURNING STRENGTH

As the MOI cannot be perfectly reconstructed from our PC basis, the skill of the linear model in forecasting $\text{MOI}_{\text{recon}}$ is not necessarily what we would like to know for practical purposes. A more relevant forecast is for MOI_{GCM} , and the bottom right panel of Fig. 3

³The linear model should not be expected to predict the variability over timescales of a few years, which is dominated by the Ekman response to atmospheric variability. This correlation increases to 0.94 if the MOI index has a 30-year filter applied, rather than the 10-year filter used throughout this study. The correlation is 0.70 using just the leading two PCs.

quantifies the skill of the linear model in predicting this decadal filtered index. It is found that for lead times less than about 4 years, damped persistence and pure persistence forecasts for MOI_{GCM} (solid grey and dashed grey lines respectively) perform best. At longer lead times the linear propagator model (solid black line) is superior, and performs better than a climatological forecast for more than 100 years. The lead time at which the linear model forecast reaches the level of 50% of the climatological variance is ~ 10 years. This compares well to three ensemble ‘perfect model’ integrations of this GCM shown by the black points and thin black lines (Collins and Sinha 2003)⁴. At zero lag, using 30 (2) PCs it is possible to estimate MOI_{GCM} with an rms error of 0.26 (0.42) Sv.

c. Predicting SST anomalies

The skill of the linear inverse model in predicting annual mean Atlantic SST anomalies is now demonstrated. Fig. 6 shows anomaly correlation maps for the LIM forecast, a damped persistence forecast, and the difference in the fraction of variance explained between the two, for 1 and 10-year lead times. The LIM shows considerable skill, even at a lead time of 10 years, over most of the Atlantic, and outperforms damped persistence in most of the basin at both lead times. The greatest increase in skill, over damped persistence, at short lead times, is in the tropical South Atlantic and in the north east Atlantic. At longer lead times, the North Atlantic Current (NAC) region shows the largest increase in skill. These results might encourage use of the LIM methodology with historical observations in the Atlantic to improve forecasts for several years to decades ahead (e.g. Newman 2007 for the Pacific Ocean), without the computational expense of running ensembles of GCM integrations.

5. Optimal perturbations

Making the forecasts as described above is important to demonstrate that the linear model has skill, but the main focus of this study is to examine the growth of anomalies. As already mentioned, none of the individual modes (ENMs) give substantial transient amplification. To obtain larger growth the damped, oscillating ENMs need to interfere constructively. This is only possible in a non-normal system as the ENMs are non-orthogonal.

a. Calculating transient amplification and optimal initial conditions

Using the definition of the norm of a vector, \mathbf{v} ,

$$\|\mathbf{v}\|_{\mathbf{N}}^2 \equiv \mathbf{v}^T \mathbf{N} \mathbf{v}, \quad (9)$$

⁴Note though that the MOC index used by Collins and Sinha (2003) is different - they used annual means of the overturning at 45°N and 670m depth, and subtracted off the Ekman part of the overturning, which is likely to enhance skill.

where \mathbf{N} is a ‘norm kernel’, we can define the amplification of our system at a lead time, τ , as,

$$A(\tau) = \frac{\mathbf{x}(\tau)^T \mathbf{N} \mathbf{x}(\tau)}{\mathbf{x}_0^T \mathbf{L} \mathbf{x}_0} = \frac{\mathbf{x}_0^T \mathbf{P}_\tau^T \mathbf{N} \mathbf{P}_\tau \mathbf{x}_0}{\mathbf{x}_0^T \mathbf{L} \mathbf{x}_0}, \quad (10)$$

where \mathbf{L} and \mathbf{N} are the initial and final norm kernels respectively. We then look for initial conditions, \mathbf{x}_0 , which maximise this amplification. This is achieved (Farrell 1988; Tziperman and Ioannou 2002) by solving the generalised eigenvalue equation,

$$\mathbf{P}_\tau^T \mathbf{N} \mathbf{P}_\tau \mathbf{x}_0 = \lambda \mathbf{L} \mathbf{x}_0, \quad (11)$$

for λ and \mathbf{x}_0 , at different lead times, τ . The largest eigenvalue, λ_{\max} , is an estimate of the largest transient amplification possible for that lead time. The optimal initial conditions are normalised so that $\mathbf{x}_0^T \mathbf{L} \mathbf{x}_0 = 1$, and the spatial patterns corresponding to \mathbf{x}_0 indicate the regions most sensitive to small perturbations, where observations would potentially be the most useful to constrain predictions.

b. Choice of norm kernel

The choice of norm kernels, \mathbf{L} and \mathbf{N} , must be considered carefully as it affects both the estimated optimal initial conditions and the growth of anomalies (Tziperman et al. 2008). Usually the two norm kernels are chosen to be the same, although there is no restriction for this to be the case. Table 1 summarises the different norms used in this study (described below), and the sensitivity to this choice is discussed later.

1) THE QUADRATIC, SST AND TOTAL VARIANCE NORMS

The simplest norm uses the identity matrix for the initial and final norm kernels ($\mathbf{L} = \mathbf{N} = \mathbf{I}$). This quadratic norm (sometimes called an energy norm) finds maximal growth in the space defined by the leading PCs. We attempt to interpret this growth by considering what any amplification means in the original data space.

It is possible to construct a norm kernel for any subset of the data,

$$\mathbf{N} = (\mathbf{W} \cdot \mathbf{E})^T (\mathbf{W} \cdot \mathbf{E}), \quad (12)$$

where \mathbf{E} is a matrix (of size $n \times p$) which consists of each of the p retained EOF patterns, restricted to the n ocean gridpoints in the field of interest, e.g., just the surface temperature data ($p = 30$ and $n = 4768$ for our dataset) to construct an SST norm (\mathbf{N}_{SST}) which we consider later, and \mathbf{W} is a matrix of weights. For the SST norm, \mathbf{W} consists of a latitude dependent area weighting. More generally, if \mathbf{E} represents the entire 3D EOF fields for temperature and salinity ($p = 30$, $n = 105150$ for our dataset) then \mathbf{N} is a norm which maximises changes in total variance. Note especially that this is *not* the identity matrix unless the various weightings (\mathbf{W}) used in the original EOF calculation (i.e. for volume, contribution to local density and local standard deviation - see HS07) are all applied to \mathbf{E}

to make it orthogonal. Thus, the quadratic norm can be interpreted as a variance norm which takes account of all these weightings. We will also consider a different variance norm (denoted by \mathbf{N}_V) which includes only the weightings for differing volumes of each grid box and for local standard deviation (although this norm still includes some covariance between temperature and salinity)⁵. To ensure comparable sizes of anomalies, \mathbf{N}_{SST} and \mathbf{N}_V are normalised to have unit determinant. When using both these norms we use the same initial and final norm kernel ($\mathbf{L} = \mathbf{N}$).

2) THE MOC NORM

Previous authors, starting with Tziperman and Ioannou (2002), have considered the transient amplification of the overturning strength by defining an MOC norm,

$$\mathbf{N}_{\text{MOC}} = \mathbf{m}\mathbf{m}^T + \varepsilon\mathbf{N}_\varepsilon, \quad (13)$$

where ε is a ‘small’ number and \mathbf{N}_ε is a suitable non-singular norm, and use \mathbf{N}_{MOC} as both the initial and final norm kernels. The extra small term (or ‘regularization’) is required as $\mathbf{m}\mathbf{m}^T$ is singular (see Tziperman and Ioannou 2002 and Tziperman et al. 2008 for more details).

This approach constrains the optimal initial MOI condition to be small ($\mathcal{O}[\varepsilon]$), resulting in a large ($\mathcal{O}[\varepsilon^{-1}]$) amplification for $\tau > 0$,

$$A_{\text{MOC}}(\tau) = \frac{\mathbf{x}_0^T \mathbf{P}_\tau^T \mathbf{N}_{\text{MOC}} \mathbf{P}_\tau \mathbf{x}_0}{\mathbf{x}_0^T \mathbf{N}_{\text{MOC}} \mathbf{x}_0} \approx \frac{\mathbf{x}_0^T \mathbf{P}_\tau^T \mathbf{m}\mathbf{m}^T \mathbf{P}_\tau \mathbf{x}_0}{\mathbf{x}_0^T \varepsilon \mathbf{N}_\varepsilon \mathbf{x}_0}, \quad (14)$$

which has the interesting feature that the initial and final norm kernels are dominated by different terms. Thus the optimal initial conditions depend strongly on the perturbing norm kernel, \mathbf{N}_ε , with the additional restriction that the initial MOI anomaly be virtually zero. If we instead choose different initial and final norms ($\mathbf{L} = \mathbf{N}_V$, and $\mathbf{N} = \mathbf{N}_{\text{MOC}}$), then we can eliminate the regularization ($\varepsilon = 0$) and find small variance anomalies which grow into large MOI anomalies, and this is the approach we adopt. This also means that the optimal initial MOI condition need not be small, i.e. it can be $\mathcal{O}[1]$.

c. Amplification under the variance norm

We first focus in detail on the maximum amplification under the variance norm. The optimal initial conditions for growth can be found by solving Eq. (11) for the leading eigenvector (\mathbf{x}_0) and maximum amplification (λ_{max}) for different norms and lead times. Fig. 7 shows the maximum amplification curves (λ_{max} as a function of τ) under the variance norm, retaining different numbers of EOFs. The maximum transient amplification increases as the

⁵Note that in Tziperman et al. (2008) the quadratic and variance norms are identical as they do not include a weighting due to a contribution to local density in their EOF estimation. We will focus on the variance norm in particular to more directly compare to Tziperman et al. (2008), but the results for the quadratic norm are very similar.

number of EOFs increases, and so the amplification found should therefore be considered as a lower bound on the possible non-normal amplification. But, Fig. 7 also demonstrates that the overall shape of the maximum amplification curve is remarkably consistent with the number of EOFs used. The weights in the leading eigenvectors, and hence the optimal initial conditions, are also fairly consistent as the number of EOFs is increased (not shown). We can therefore be reasonably confident that the optimal initial conditions and amplification process are robust.

It is found that the maximum transient amplification, using 30 EOFs, occurs at 36 years and involves an amplification factor of ≈ 10 . The thin black curve in Fig. 7 shows the evolution of this largest growing mode. The level of maximum amplification is slightly larger than that found for the GFDL CM2.1 model by Tziperman et al. (2008), although the time to reach maximum amplification is far longer in HadCM3 (36 years compared to ~ 8 years). We now focus on the mechanisms of amplification of this maximal growth.

The evolution of temperature and salinity anomalies from the initial (\mathbf{x}_0) to final ($\mathbf{P}_\tau \mathbf{x}_0$) states, projected into physical space and integrated from the surface to 1800m depth, is shown in Fig. 8. The initial state has fairly localised, and shallow (Fig. 9) anomalies concentrated in the North Atlantic and Nordic Seas, which grow into anomalies affecting most of the Atlantic basin, with a strong north-east/south-west dipole structure. The inference is that a small anomaly with the structure of the initial state will grow into the final state over 36 years. It is of interest that the anomalies tend to grow in volume rather than amplitude and are substantially density compensating, and therefore involve changes in the spiciness of the water masses. The evolution of these anomalies shows propagation from the North Atlantic southward along the western boundary into the tropical Atlantic, probably due to a combination of wave and advective processes, followed by spreading into the interior.

Fig. 9 shows how these anomalies are distributed with depth, integrated over two latitude bands. The dashed lines represent the initial state and the solid lines are the final state. It can be seen that the initial anomalies are concentrated in the top 500m, especially in the sub-polar gyre. The final state is dominated by deeper anomalies (below 500m) in the tropical Atlantic. The locations highlighted in the initial state indicate where ‘observations’ would be most valuable for constraining predictions. The shallow nature of the initial anomalies is encouraging, as these may be more easily observable than anomalies at deeper levels. Also note that although the sign of the anomalies is arbitrary in this linear model, the GCM behaves similarly for both signs of anomalies (see Fig. 12 later).

Tziperman et al. (2008) also found that the optimal initial conditions were located in the far North Atlantic in the GFDL CM2.1 model. In their maps of dynamic topography there are also hints of a similar western boundary propagation. It is encouraging to see similar processes at work in two different models.

The evolution of the MOC during this amplification is also of interest. The solid black line in Fig. 10 shows that the negative $\text{MOI}_{\text{recon}}$ anomaly amplifies for the first ten years, and then recovers. Thorpe et al. (2001) showed that the MOC strength in HadCM3 is strongly correlated with the meridional density gradient in the North Atlantic, and the

evolution of anomalous meridional density gradients in Fig. 8 shows the same effect with anomalous low density in the far North Atlantic occurring when the MOC is anomalously weak. De Coëtlogon et al. (2006) also found correlations between Gulf Stream position and transport changes and the MOC in several GCMs, and similar relationships between Gulf Stream transport and gradients in potential energy anomalies have been identified in observations (Curry and McCartney 2001).

Figs. 8 and 10 indicate that there may be two stages to the amplification. A faster (~ 10 year) growth due to changes in the high latitudes - the optimal perturbation is a low density anomaly (dominated by low salinity) in the surface layers of the Nordic Seas, which inhibits convection. The mean density of water flowing through Denmark Strait is also lowered (not shown), and these factors are known to lead decreases in the strength of the MOC in HadCM3 (Hawkins and Sutton 2008). There is also propagation along the western boundary and a corresponding change in the MOC. After a few years, the convection can restart and the MOC recovers. During the second stage (years ~ 10 onwards), there is slower growth spreading the warm and salty anomalies to greater depths and into the ocean interior, especially in the tropics.

d. Comparing different norms

One of the uses for optimal initial conditions is to perturb ocean analyses to help design efficient decadal climate forecast ensembles (e.g. Smith et al. 2007; Meehl et al. 2008). Therefore we now examine the sensitivity of our results to the choice of norm for a fixed optimisation lead time of 10 years. Fig. 11 shows the optimal initial conditions for each norm. The patterns are very similar, with large values found in the far North Atlantic, especially in the Labrador Seas and Nordic Seas - i.e. the regions of deep ocean convection. The pattern for the SST norm is slightly different with more weight in the tropical North Atlantic, especially in salinity. However, after 10 years growth, the amplified patterns are fairly similar (not shown).

Fig. 10 considers the development of MOI anomalies for the different norms for the fixed optimisation time of 10 years. The evolution is similar for the variance and MOC norms, and also the quadratic norm (not shown), with growth to the largest MOC extremum⁶ occurring at $\sim 8 - 12$ years, followed by a decay. The SST norm shows a similar evolution but with the largest MOC anomaly at around 20 years.

Table 1 gives the lead time of largest possible amplification for each norm. Using the SST norm there is larger non-normal amplification - a factor of ~ 32 at a lead time of 5 years. The MOC norm shows a far smaller maximal amplification of just ~ 1.35 from 1 to 10 years.

These results suggest that, (i) the optimal initial condition is dependent on the norm chosen, (ii) the growth of anomalies over the first decade appears dominated by an MOC perturbation under all norms, and (iii) the timescale to maximum amplification, and the pattern at maximum amplification are sensitive to the norm chosen.

⁶Note again that the sign of the anomaly is arbitrary.

e. Does this amplification actually occur?

The calculation of these optimal initial conditions is a mathematical construct, and so it is important to examine whether the amplification described actually occurs. One way to test this would be to impose the optimal initial conditions as anomalies in the full GCM and run an ensemble of integrations. Firstly, this is a computationally expensive test of the kind we are trying to avoid, and secondly, these patterns are ‘average’ amplifying perturbations, which are not optimised for any particular initial condition. A simpler, and perhaps more relevant test, is to use the natural internal variability generated by the GCM.

Although it is extremely unlikely that, at any one particular time, the anomalies in the GCM control integration will exactly match the optimal initial condition estimated above, we can study the projection of the GCM anomalies, for all 1100 years, onto the optimal initial and final states (see PS95).

This simple estimate of the relevance of the optimal initial condition is shown in Fig. 12 which plots, in the space of the leading PCs, the projection of the state vector, $\mathbf{x}(t)$, onto the optimal initial condition, $\mathbf{x}_0(\tau)$, against the projection of $\mathbf{x}(t + \tau)$ onto the maximally amplified state, $\mathbf{P}_\tau \mathbf{x}_0(\tau)$, for all t , for a fixed optimisation time of $\tau = 10$ years, for each norm. If this linear non-normal amplification is occurring in the GCM then all the points should lie near the dashed lines (a slope of unity). The relation will not be perfect due to the integrated effect of noise throughout the evolution of the anomaly, but it is reassuring that the best-fit slopes are positive, with good correlations. This demonstrates that this type of amplification does indeed occur in the GCM and that our linear model is capturing this. Although there are small differences from a slope of unity, departures from linearity are modest.

Although the structure of the growth appears to occur in the same way as the linear model suggests, the observed level of amplification in the GCM typically reaches $\sim 50\%$ of the maximum linear estimate (not shown). This is not surprising as we expect the non-linearities in the GCM to impact our linear estimates.

These results give us confidence that our estimated optimal initial conditions are realistic and giving rise to amplification in the GCM in the way described by our simple linear model.

6. Conclusions and implications

We have analysed the decadal variability and predictability of 3D Atlantic Ocean anomalies in a GCM (HadCM3) using a linear inverse modelling approach. The main findings can be summarised as follows:

- the evolution of temperature and salinity in the Atlantic Ocean as captured by the leading 30 three-dimensional EOFs can be effectively described by a linear dynamical system, forced by white noise.
- forecasts, using the linear model, of Atlantic SSTs, the strength of the overturning,

and of the leading PCs are more skillful than damped persistence and climatological forecasts for several decades.

- significant non-normal amplification by a factor of ~ 10 can occur in the linear model framework under the variance norm over 36 years. Similar processes of amplification are also found in the natural variability of the full GCM.
- the regions from which growth occurs under the variance, quadratic, and MOC norms are found in the top 500m of the far North Atlantic. Anomalies in this region amplify to fill most of the Atlantic basin, especially at depth, over 1-3 decades. The mechanism of amplification appears to initially involve anomalies in the Nordic Seas, followed by southward propagation of anomalies along the western boundary affecting the MOC, and subsequent spreading into the interior.
- the structure of the optimal initial conditions is sensitive to the choice of norm, whereas the initial growth appears similar and dominated by MOC-related basin scale changes for all norms considered.

Tziperman et al. (2008) used a similar methodology to that described here, using a quadratic norm, on the GFDL CM2.1 model and found a transient amplification of comparable amplitude which occurred over the first 10 years - similar to the first stage of growth found here for HadCM3. They concluded that the predictability in this model decays rapidly over 10 years, and this is also seen in the errors in their forecasts, which grow more rapidly than for HadCM3. In HadCM3 the transient growth is slower, indicating greater predictability in this model. This is also demonstrated by the significant skill found in the forecasts, even at long lead times, and also by previous studies (e.g. Collins and Sinha 2003).

There are two main reasons for performing this kind of analysis. Firstly, the initial perturbations indicate where observations would be most valuable to help constrain decadal climate predictions. The similar location (far North Atlantic) for the optimal initial conditions found in both HadCM3 and GFDL CM2.1 is encouraging, and suggests additional ocean observations in these regions could be valuable, though the spatial density of ocean observations that would be required is still an open question. Secondly, any future operational decadal climate prediction system will rely on ensemble forecasts from analysed initial conditions, and will require optimal perturbations, such as these, to efficiently explore the uncertainty in the initial state, and hence provide more accurate forecast uncertainties.

It should be noted that, although we have attempted to justify the various assumptions about the linearity of the system, it needs to be rigorously demonstrated that the identified regions are indeed the most sensitive to perturbations using integrations of the full GCM (e.g. Kleeman et al. 2003). A further paper will explore this issue in more detail. Further work will also investigate the application of these methods to different GCMs.

Acknowledgments.

We thank Myles Allen for suggesting the use of the linear inverse modelling approach, and Cécile Penland for valuable discussions. We also thank Laure Zanna and the anonymous reviewer for their comments which helped improve the paper. EH is funded by the UK Natural Environment Research Council under the thematic Rapid Climate Change programme. RS is supported by a Royal Society University Research Fellowship.

APPENDIX A

Methodology details

Constructing a linear propagator

Here we give more details (also see Penland and Sardeshmukh 1995) of the estimation of \mathbf{B} , starting from Eq. (2). If we assume (for now) that the forcing term (ξ) is zero then Eq. (2) is a purely linear system, and we can introduce a linear propagator, \mathbf{G} , so that,

$$\mathbf{x}(t + \delta) = \mathbf{G}(\delta)\mathbf{x}(t) \quad (\text{A1})$$

where \mathbf{x} is our PC state vector and δ is the lead time. This assumption of linearity is not going to be strictly true, and so we assume that the non-linear terms can be considered as a forcing term (ξ) in Eq. (2). This forcing is generally taken to be Gaussian white noise (in time). These assumptions of linearity and Gaussianity are tested for our system in Appendix B. The optimal linear propagator $\mathbf{G}(\delta)$ can be found from the covariance matrices of \mathbf{x} (PS95),

$$\mathbf{G}(\delta) = \frac{\mathbf{C}(\delta)}{\mathbf{C}(0)} \quad (\text{A2})$$

where

$$\mathbf{C}(\delta) = \langle \mathbf{x}(t + \delta) \mathbf{x}^T(t) \rangle, \quad (\text{A3})$$

$$\mathbf{C}(0) = \langle \mathbf{x}(t) \mathbf{x}^T(t) \rangle, \quad (\text{A4})$$

where $\langle \cdot \rangle$ denotes an average over all t . $\mathbf{C}(0)$ is a diagonal matrix with elements proportional to the fraction of the variance explained by each EOF. In this notation,

$$\mathbf{B} = \frac{\ln[\mathbf{G}(\delta)]}{\delta}, \quad (\text{A5})$$

and should be independent of the lead time (δ) used to estimate it. The general linear propagator is then (c.f. Eqn. 4),

$$\mathbf{P}_\tau = \exp\left(\frac{\tau}{\delta} \ln[\mathbf{G}(\delta)]\right) = \exp(\tau\mathbf{B}). \quad (\text{A6})$$

Although, in principle, we are free to choose any lead time, δ , to calculate \mathbf{B} there are physical constraints which limit this choice. To ensure that the modes of variability are physically realistic decaying modes it is necessary that the eigenvalues of $\mathbf{G}(\delta)$ have positive real part and magnitude less than unity. It is found for our 3D EOF system that using a lead time of $\delta = 1$ year will satisfy this constraint, but that longer lead times do not. All the results presented here therefore use the propagator derived for $\delta = 1$ year.

Empirical Normal Modes

The eigenvectors of \mathbf{B} are sometimes known as the Empirical Normal Modes (ENMs, PS95) and the corresponding eigenvalues, β_j , are generally complex,

$$\beta_j = \sigma_j + i\omega_j. \quad (\text{A7})$$

Table 2 summarises the properties of the ENMs and shows that two modes in our system are exponentially decaying modes ($\omega_j = 0$) and the rest are oscillatory, complex-conjugate pairs with decay times $-1/\sigma_j$, and periods $2\pi/\omega_j$. The shortest period for one of these oscillatory modes is 4.4 years (ENM 3/4). Erroneous results for \mathbf{G} can be obtained if the lag (δ) is chosen to be larger than half the shortest period, and large uncertainties can occur near this half period (PS95). This is the so-called ‘Nyquist problem’, and is the reason why the propagator matrix could not be reliably estimated for $\delta > 1$ year in our 30 EOF system.

Noise EOFs

As the individual ENMs decay there must be an energy input from the stochastic forcing (ξ) to ensure stationary statistics (a fluctuation-dissipation relation). This relation (Penland and Matrosova 1994; PS95) is,

$$\mathbf{B}\mathbf{C}(0) + \mathbf{C}(0)\mathbf{B}^T + \mathbf{Q} = 0, \quad (\text{A8})$$

and defines the covariance matrix of the forcing, $\mathbf{Q} = \langle \xi \xi^T \rangle dt$. Although the forcing is white noise in time, the eigenvectors of \mathbf{Q} correspond to coherent, three-dimensional spatial patterns of forcing, sometimes called ‘noise EOFs’. The most likely source for noise in the ocean is the atmosphere and so it is hoped that these noise EOFs are dominated by surface anomalies which resemble known surface flux patterns due to coherent atmospheric forcing states. In our case the leading noise EOFs are dominated by anomalies in the top $\sim 350\text{m}$, with the spatial pattern of the first noise EOF (representing around 13% of the variance) is similar to the observed ocean response to changes in the North Atlantic Oscillation (not shown).

APPENDIX B

Validation of the assumptions in the methodology

There are several tests which can be used to justify the assumptions made in constructing the linear model, and to check that our system is indeed acting like a Gaussian white noise forced version of Eq. (2). Section 4a analyses the skill of the forecasts produced, and here we consider some further tests.

Are the linearity assumptions valid?

The evolution of the PCs will not be perfectly linear, but the skill of the forecasts shown in Fig. 3 give us confidence in this assumption. Also, as previously discussed, for a linear system, \mathbf{B} should be independent of δ . Unfortunately we cannot check this as \mathbf{B} can only be reliably calculated for $\delta = 1$ year, using 30 EOFs. If fewer EOFs are used then this test is reasonably well passed (not shown).

Is the noise forcing consistent?

As the noise matrix, \mathbf{Q} , is a covariance matrix it should be positive definite, i.e. the eigenvalues should all be positive. It is found that, when using 30 EOFs, two of the eigenvalues are slightly negative ($\approx -10^{-5}$). We also find that if only the leading 15 EOFs are retained then all the eigenvalues are positive. This may indicate the presence of a small amount of non-linearity found in the low-variance, higher EOFs, or that we do not have enough data to reliably determine the higher EOFs. Previous studies (e.g PS95) have found similar results and concluded that it does not influence their results significantly.

Are the statistics consistent?

Another approach to testing the linear model assumptions is to integrate a forward model, constructed from Eq. (2), and generate artificial data with the same linear propagator as that estimated from the GCM data (see Penland and Matrosova 1994 and PS95 for details). This artificial data should have similar statistics to the real data. Newman (2007) compared the power spectra of the components output by the forward model with the power spectra derived directly from the GCM data and from observations and we perform a similar test here.

We integrated our forward model for 55 000 years, with a time step of 1 month, and split the output into 50 independent segments, so that each segment had the same length (1100 years) as the GCM data, to act as a Monte Carlo ensemble. The stochastic forcing used in the forward model is generated using the noise matrix, \mathbf{Q} , with the two negative eigenvalues suppressed.

The solid lines in Fig. 13 show the power spectra for the leading five PCs and these should be compared to the dashed lines which indicate the ensemble mean from the forward model, with 95% confidence limits shaded grey. It can be seen that the power spectra derived directly from the data are within the estimated confidence levels for the leading five principal components, indicating that the linear model can reproduce the frequency statistics of the distribution well. The large uncertainty range in the power spectra indicates that even the statistics of 1100 year segments can be very variable and each segment may not fully represent all the frequencies of natural variability. Fig. 14 shows the same for the domain means of SST, SSS, temperature and salinity at a depth of 1500m, and the MOI. It is not surprising to see that the linear model is better at representing the variability of the deep ocean than the surface layers. The lower panel of Fig. 14 indicates that not all of the MOI variability is completely represented by the linear model. Also seen is the similarity of the spectra for MOI_{GCM} and $\text{MOI}_{\text{recon}}$ for periods greater than about 20 years, except for the 25 year spectral peak which is missing in $\text{MOI}_{\text{recon}}$. These differences will limit the skill of the linear model to predict the overturning indices.

Considering Gaussianity

If a Gaussian white noise forced linear model is a good representation of our system then the statistics of \mathbf{x} should be Gaussian, but it is found that (only) the leading two PCs have significantly non-Gaussian skewness. This could be an indication of non-linear processes which are not well represented by Gaussian white noise, or due to the low frequency variability (Fig. 13) found in these two leading PCs which is not fully sampled in the 1100 years of analysed data. As a further check we examined the statistics of different segments of the stochastically forced forward model described above and found significantly non-Gaussian skewness for the leading two PCs in many segments. The statistics were only Gaussian if the segment considered was long enough to capture enough cycles of the coherent oscillation. Even so, the skewness of the real PC1 is outside the 95% confidence limits derived from the forward model (not shown). Non-Gaussianity is possible if the noise is multiplicative rather than additive (e.g. Sura and Sardeshmukh 2008), but this is a small caveat on our results.

REFERENCES

- Alexander, M. A., L. Matrosova, C. Penland, J. D. Scott, and P. Chang, 2008: Forecasting Pacific SSTs: Linear Inverse Model predictions of the PDO. *J. Climate*, **21**, 385–402, doi:10.1175/2007JCLI1849.1.
- Collins, M. and B. Sinha, 2003: Predictability of decadal variations in the thermohaline circulation and climate. *Geophys. Res. Lett.*, **30**, 1306, doi:10.1029/2002GL016504.
- Collins, M., et al., 2006: Interannual to decadal climate predictability in the North Atlantic: a multimodel-ensemble study. *J. Climate*, **19**, 1195–1202, doi:10.1175/JCLI3654.1.
- Curry, R. G. and M. S. McCartney, 2001: Ocean gyre circulation changes associated with the North Atlantic Oscillation. *J. Phys. Ocean.*, **31**, 3374–3400.
- De Coëtlogon, G., C. Frankignoul, M. Bentsen, C. Delon, H. Haak, S. Masina, and A. Par-daens, 2006: Gulf Stream variability in five oceanic general circulation models. *J. Phys. Ocean.*, **36**, 2119–2135.
- Dong, B. and R. T. Sutton, 2005: Mechanism of interdecadal thermohaline circulation variability in a coupled ocean-atmosphere GCM. *J. Climate*, **18**, 1728.
- Farrell, B., 1988: Optimal excitation of neutral Rossby waves. *J. Atmos. Sci.*, **45**, 163 – 172.
- Farrell, B. and P. J. Ioannou, 2001: Accurate low-dimensional approximation of the linear dynamics of fluid flow. *J. Atmos. Sci.*, **58**, 2771–2789.
- Gordon, C., C. Cooper, C. A. Senior, H. Banks, J. M. Gregory, T. C. Johns, J. F. B. Mitchell, and R. A. Wood, 2000: The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre coupled model without flux adjustments. *Climate Dyn.*, **16**, 147–168.
- Griffies, S. M. and K. Bryan, 1997: Predictability of North Atlantic multidecadal climate variability. *Science*, **275**, 181–184.
- Hawkins, E. and R. Sutton, 2007: Variability of the Atlantic thermohaline circulation described by three-dimensional empirical orthogonal functions. *Climate Dyn.*, **29**, 745–762, doi:10.1007/s00382-007-0263-8.
- Hawkins, E. and R. Sutton, 2008: Potential predictability of rapid changes in the Atlantic meridional overturning circulation. *Geophys. Res. Lett.*, **35**, L11 603, doi:10.1029/2008GL034059.

- Keenlyside, N. S., M. Latif, J. Jungclauss, L. Kornbluh, and E. Roeckner, 2008: Advancing decadal-scale climate prediction in the North Atlantic sector. *Nature*, **453**, 84–88, doi:10.1038/nature06921.
- Kleeman, R., Y. Tang, and A. M. Moore, 2003: The calculation of climatically relevant singular vectors in the presence of weather noise as applied to the ENSO problem. *J. Atmos. Sci.*, **60**, 2856 – 2868.
- Lohmann, G. and J. Schneider, 1999: Dynamics and predictability of Stommel’s box model: a phase space perspective with implications for decadal climate variability. *Tellus A*, **51**, 326–336.
- Meehl, G. A., et al., 2008: Decadal prediction: can it be skillful? *Bull. Amer. Met. Soc.*, submitted.
- Molteni, F., R. Buizza, T. N. Palmer, and T. Petroliagis, 1996: The ECMWF Ensemble Prediction System: methodology and validation. *Q. J. Royal Met. Soc.*, **122**, 73–119, doi:10.1002/qj.49712252905.
- Newman, M., 2007: Interannual to decadal predictability of tropical and North Pacific sea surface temperatures. *J. Climate*, **20**, 2333 – 2356, doi:10.1175/JCLI4165.1.
- Pardaens, A. K., H. T. Banks, J. M. Gregory, and P. R. Rowntree, 2003: Freshwater transports in HadCM3. *Climate Dyn.*, **21**, 177–195, doi:10.1007/s0038200303246.
- Penland, C., 1989: Random forcing and forecasting using principal oscillation pattern analysis. *Mon. Wea. Rev.*, **117**, 2165 – 2185.
- Penland, C., 1996: A stochastic model of IndoPacific sea surface temperature anomalies. *Physica D*, **98**, 534–558.
- Penland, C. and L. Matrosova, 1994: A balance condition for stochastic numerical models with application to the El Niño-Southern Oscillation. *J. Climate*, **7**, 1352 – 1372.
- Penland, C. and P. D. Sardeshmukh, 1995: The optimal growth of tropical sea surface temperature anomalies. *J. Climate*, **8**, 1999 – 2024.
- Pohlmann, H., M. Botzet, M. Latif, A. Roesch, M. Wild, and P. Tschuck, 2004: Estimating the decadal predictability of a coupled AOGCM. *J. Climate*, **17**, 4463–4472.
- Sevellec, F., T. Huck, M. B. Jelloul, N. Grima, J. Vialard, and A. Weaver, 2008: Optimal surface salinity perturbations of the meridional overturning and heat transport in a global ocean general circulation model. *J. Phys. Ocean.*, **38**, 2739–2754.
- Smith, D. M., S. Cusack, A. W. Colman, C. K. Folland, G. R. Harris, and J. M. Murphy, 2007: Improved surface temperature prediction for the coming decade from a global climate model. *Science*, **317**, 796–799, doi:10.1126/science.1139540.

- Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. M. B. Tignor, and H. L. Miller, (Eds.) , 2007: *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK.
- Sura, P. and P. Sardeshmukh, 2008: A global view of non-Gaussian SST variability. *J. Phys. Ocean.*, **38**, 639–647, doi:10.1175/2007JPO3761.1.
- Sutton, R. T. and D. L. R. Hodson, 2005: Atlantic Ocean forcing of North American and European summer climate. *Science*, **309**, 115–118, doi:10.1126/science.1109496.
- Thorpe, R. B., J. M. Gregory, T. C. Johns, R. A. Wood, and J. F. B. Mitchell, 2001: Mechanisms determining the Atlantic thermohaline circulation response to greenhouse gas forcing in a non-flux-adjusted coupled climate model. *J. Climate*, **14**, 3102–3116.
- Tziperman, E. and P. J. Ioannou, 2002: Transient growth and optimal excitation of thermohaline variability. *J. Phys. Ocean.*, **32**, 3427–3435.
- Tziperman, E., L. Zanna, and C. Penland, 2008: Non-normal thermohaline circulation dynamics in a coupled ocean-atmosphere GCM. *J. Phys. Ocean.*, **38**, 588 – 604, doi: 10.1175/2007JPO3769.1.
- Vellinga, M. and P. Wu, 2004: Low-latitude freshwater influence on centennial variability of the Atlantic thermohaline circulation. *J. Climate*, **17**, 4498–4511.
- Zanna, L. and E. Tziperman, 2005: Non-normal amplification of the thermohaline circulation. *J. Phys. Ocean.*, **35**, 1593–1605.

List of Figures

- 1 Left hand axis scale: the eigenspectrum of the leading 30 EOFs (black line, circles). Right hand axis scale: cumulative variance explained (dark grey line, squares) and cumulative variance of the MOI explained (light grey line, diamonds). 25
- 2 The leading four PCs (grey lines) and forecasts using the linear propagator, \mathbf{P}_τ (black lines), for lead times of up to 50 years, plotted at 50 year intervals. The linear propagator seems to show significant skill, even at long lead times. 26
- 3 Root mean square (rms) errors, relative to a climatological forecast, for forecasts of quantities as labelled. Solid black lines: using the linear propagator, \mathbf{P}_τ . Dashed black lines: ‘perfect’ linear model (Eq. (7) for PCs, Eq. (8) for $\text{MOI}_{\text{recon}}$). Solid grey lines: a damped persistence forecast. Dashed grey lines: a persistence forecast. A climatological forecast (i.e. always predicting zero anomaly) is indicated by the thin horizontal line equal to 1. The thin horizontal dashed line is the level at which the error variance reaches 50% of the climatological variance, an approximate level of significant predictability (Griffies and Bryan 1997). In the bottom right panel the solid black line does not start at zero as MOI_{GCM} is not perfectly reconstructed from our PC basis. There is no perfect linear model for this forecast. The small black points in this panel show results from Collins and Sinha (2003) using ‘perfect model’ experiments. 27
- 4 Root mean square (rms) errors, relative to a climatological forecast, for forecasts of quantities as labelled. Solid black lines: using the linear propagator, \mathbf{P}_τ , for forecasts of all 1100 years, without cross-validation. Dashed black lines: using the linear propagator, \mathbf{P}_τ , for forecasts of just the last 100 years. Solid grey lines: using a linear propagator estimated using just the first 1000 years, for forecasts of the independent last 100 years. A climatological forecast (i.e. always predicting zero anomaly) is indicated by the thin horizontal line equal to 1. The thin horizontal dashed line is the level at which the error variance reaches 50% of the climatological variance, an approximate level of significant predictability (Griffies and Bryan 1997). 28
- 5 Overturning strength: The ‘true’ MOI index (MOI_{GCM} , dashed grey line) is shown with the reconstructed MOI index (solid grey line). The correlation between these two quantities is 0.90. Forecasts using the linear propagator, \mathbf{P}_τ (black lines), are shown for lead times of up to 50 years, plotted at 50 year intervals. 29
- 6 Forecast skill, shown by anomaly correlations of annual mean Atlantic SSTs from the forecasts and from the GCM, at lead times of 1 year (left column) and 10 years (right column) for the LIM (top row), and damped persistence (middle row) forecasts. The bottom row shows the difference in percentage of variance explained between the two forecasts. The linear inverse model beats damped persistence over most of the Atlantic, and shows significant skill at a lead time of 10 years. The grey contours in the bottom left panel show the regions where damped persistence beats the LIM by more than 15%. 30

7	The maximum amplification curves for the variance norm, retaining various numbers of EOFs as labelled. The largest transient amplification, using 30 EOFs, is a factor of ~ 10 , at a lead time of 36 years. The thin black line shows the evolution of this maximally amplifying mode.	31
8	The evolution of the maximally amplifying eigenvector under the variance norm, integrated from the surface to a depth of 1800m. Left column: temperature (in K). Middle column: salinity (in psu), multiplied by five. Right column: density (in kg m^{-3}), multiplied by seven. Top row: initial state. Second row: evolved state after 10 years. Third row: evolved state after 22 years. Bottom row: maximally amplified state after 36 years. The colour scale is the same in every panel and is arbitrary. White regions represent small anomalies of either sign.	32
9	Depth profiles of the maximally amplifying eigenvector under the variance norm, for temperature and salinity integrated over latitude bands as labelled. Dashed lines are the optimal initial perturbation, and the solid lines are the final state (after 36 years). The scales are arbitrary representations in the correct units.	33
10	The evolution of $\text{MOI}_{\text{recon}}$ during non-normal amplification. Each curve shows the temporal evolution of the maximally amplifying mode for the labelled norm, with maximum amplification times as indicated.	34
11	The initial conditions for maximal amplification over 10 years for each norm, integrated from the surface to a depth of 1800m. Left column: temperature (in K). Right column: salinity (in psu), multiplied by five. The colour scale is the same in every panel and is arbitrary. White regions represent small anomalies of either sign.	35
12	Scatter plots of the projection of the state vector, $\mathbf{x}(t)$, onto the optimal initial condition, $\mathbf{x}_0(\tau)$, against the projection of $\mathbf{x}(t + \tau)$ onto the maximally amplified state, $\mathbf{P}_\tau \mathbf{x}_0(\tau)$, for all t , and fixed $\tau = 10$ years. The dashed line in each panel has a slope of unity. (a) Using the variance norm. (b) Using the quadratic norm. (c) Using the SST norm. (d) Using the MOC norm.	36
13	Power spectra for the leading five principal components. Solid lines: for 1100 years of the GCM data, \mathbf{x} . Dashed lines: ensemble mean of 50 separate segments of 1100 years from a forward model with 95% confidence limits shaded grey.	37
14	Power spectra for SST, SSS, temperature and salinity at a depth of 1500m, and the MOI. Solid lines: for 1100 years of the GCM data, \mathbf{x} . Dashed lines: ensemble mean of 50 separate segments of 1100 years from a forward model with 95% confidence limits shaded grey. In the MOI panel the thick solid line represents $\text{MOI}_{\text{recon}}$ and the thin solid line represents MOI_{GCM}	38

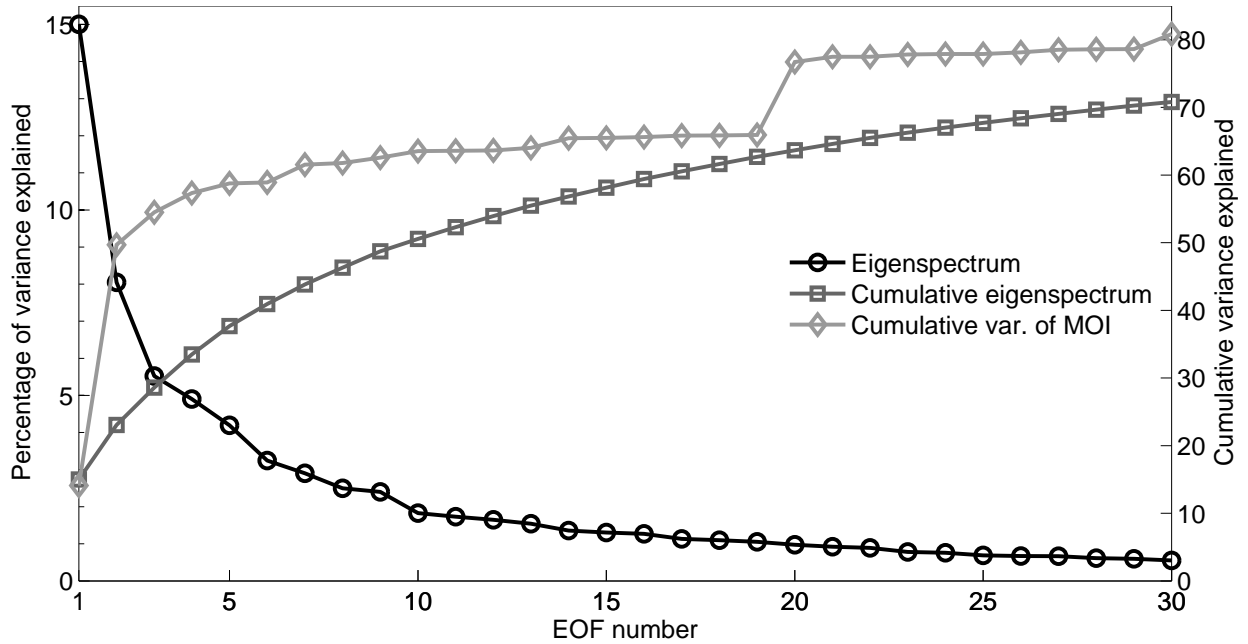


FIG. 1. Left hand axis scale: the eigenspectrum of the leading 30 EOFs (black line, circles). Right hand axis scale: cumulative variance explained (dark grey line, squares) and cumulative variance of the MOI explained (light grey line, diamonds).




FIG. 2. The leading four PCs (grey lines) and forecasts using the linear propagator, \mathbf{P}_τ (black lines), for lead times of up to 50 years, plotted at 50 year intervals. The linear propagator seems to show significant skill, even at long lead times.




FIG. 3. Root mean square (rms) errors, relative to a climatological forecast, for forecasts of quantities as labelled. Solid black lines: using the linear propagator, \mathbf{P}_τ . Dashed black lines: ‘perfect’ linear model (Eq. (7) for PCs, Eq. (8) for $\text{MOI}_{\text{recon}}$). Solid grey lines: a damped persistence forecast. Dashed grey lines: a persistence forecast. A climatological forecast (i.e. always predicting zero anomaly) is indicated by the thin horizontal line equal to 1. The thin horizontal dashed line is the level at which the error variance reaches 50% of the climatological variance, an approximate level of significant predictability (Griffies and Bryan 1997). In the bottom right panel the solid black line does not start at zero as MOI_{GCM} is not perfectly reconstructed from our PC basis. There is no perfect linear model for this forecast. The small black points in this panel show results from Collins and Sinha (2003) using ‘perfect model’ experiments.




FIG. 4. Root mean square (rms) errors, relative to a climatological forecast, for forecasts of quantities as labelled. Solid black lines: using the linear propagator, \mathbf{P}_τ , for forecasts of all 1100 years, without cross-validation. Dashed black lines: using the linear propagator, \mathbf{P}_τ , for forecasts of just the last 100 years. Solid grey lines: using a linear propagator estimated using just the first 1000 years, for forecasts of the independent last 100 years. A climatological forecast (i.e. always predicting zero anomaly) is indicated by the thin horizontal line equal to 1. The thin horizontal dashed line is the level at which the error variance reaches 50% of the climatological variance, an approximate level of significant predictability (Griffies and Bryan 1997).




FIG. 5. Overturning strength: The ‘true’ MOI index (MOI_{GCM} , dashed grey line) is shown with the reconstructed MOI index (solid grey line). The correlation between these two quantities is 0.90. Forecasts using the linear propagator, \mathbf{P}_τ (black lines), are shown for lead times of up to 50 years, plotted at 50 year intervals.

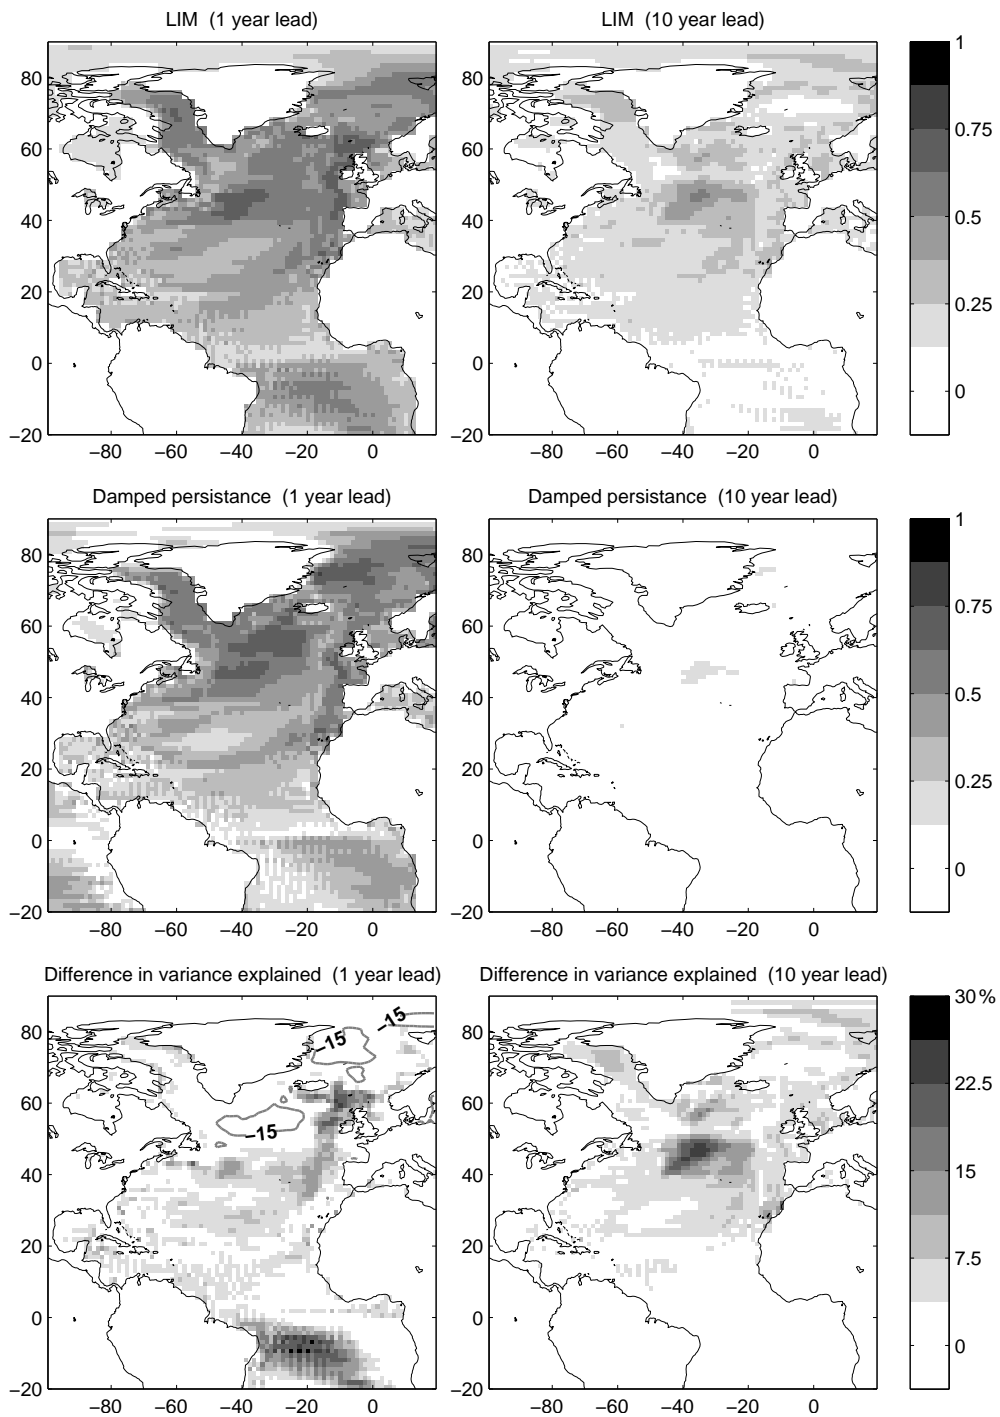


FIG. 6. Forecast skill, shown by anomaly correlations of annual mean Atlantic SSTs from the forecasts and from the GCM, at lead times of 1 year (left column) and 10 years (right column) for the LIM (top row), and damped persistence (middle row) forecasts. The bottom row shows the difference in percentage of variance explained between the two forecasts. The linear inverse model beats damped persistence over most of the Atlantic, and shows significant skill at a lead time of 10 years. The grey contours in the bottom left panel show the regions where damped persistence beats the LIM by more than 15%.

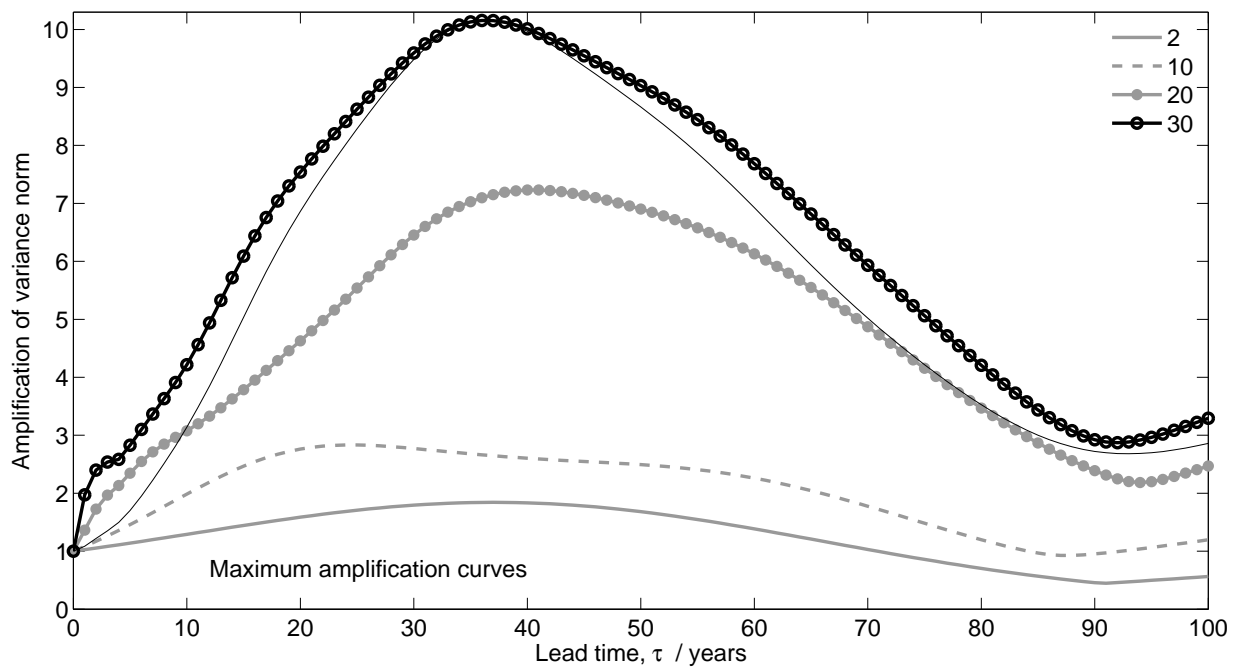


FIG. 7. The maximum amplification curves for the variance norm, retaining various numbers of EOFs as labelled. The largest transient amplification, using 30 EOFs, is a factor of ~ 10 , at a lead time of 36 years. The thin black line shows the evolution of this maximally amplifying mode.

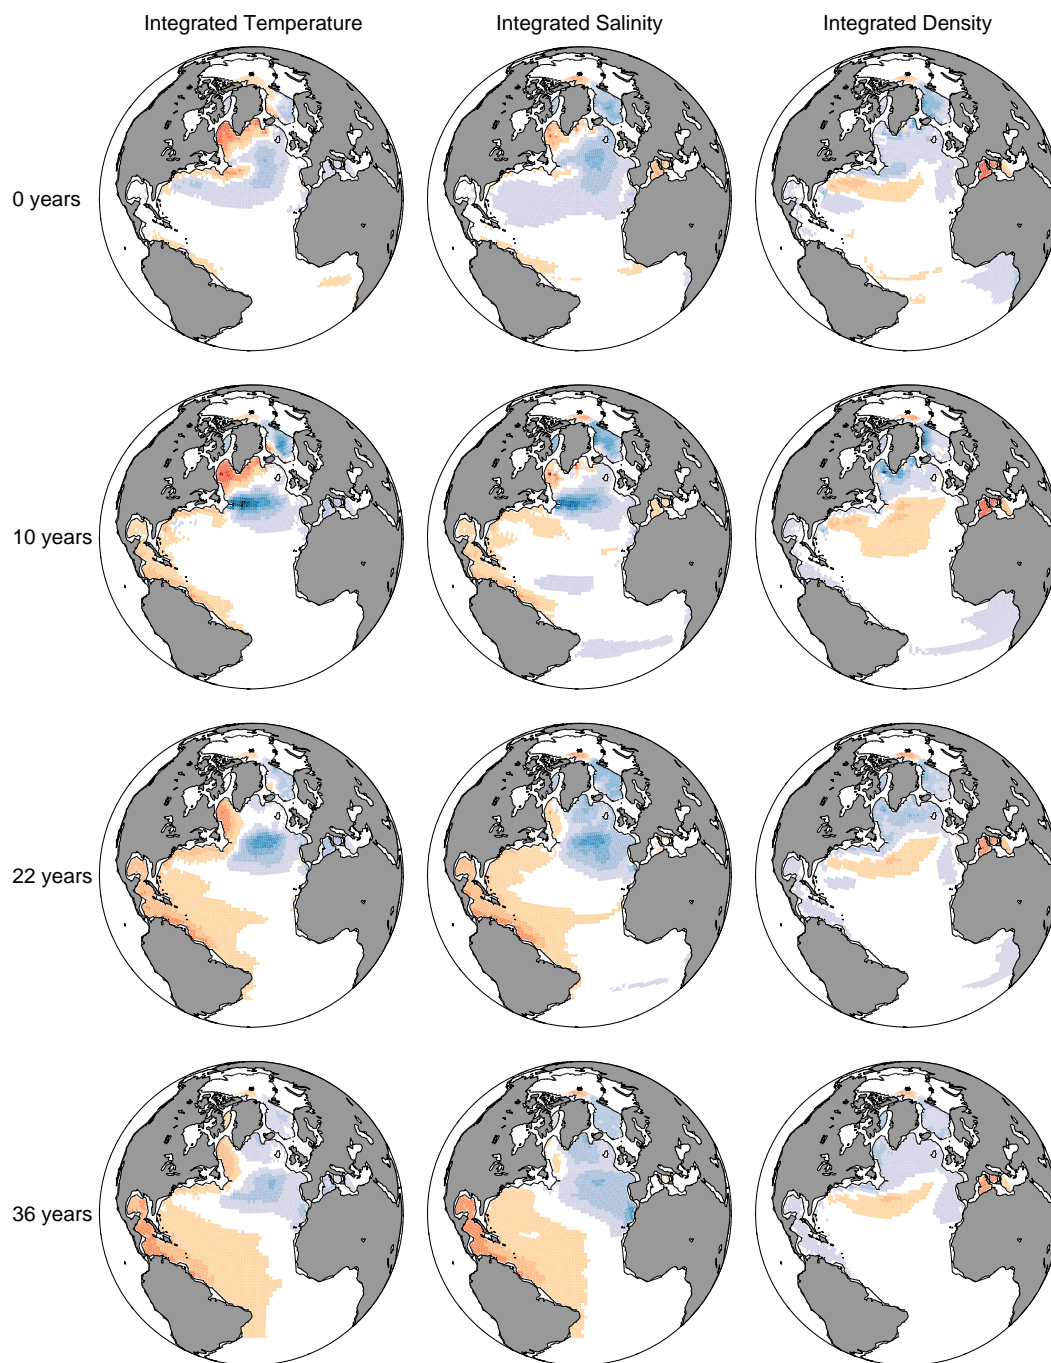


FIG. 8. The evolution of the maximally amplifying eigenvector under the variance norm, integrated from the surface to a depth of 1800m. Left column: temperature (in K). Middle column: salinity (in psu), multiplied by five. Right column: density (in kg m^{-3}), multiplied by seven. Top row: initial state. Second row: evolved state after 10 years. Third row: evolved state after 22 years. Bottom row: maximally amplified state after 36 years. The colour scale is the same in every panel and is arbitrary. White regions represent small anomalies of either sign.




FIG. 9. Depth profiles of the maximally amplifying eigenvector under the variance norm, for temperature and salinity integrated over latitude bands as labelled. Dashed lines are the optimal initial perturbation, and the solid lines are the final state (after 36 years). The scales are arbitrary representations in the correct units.




FIG. 10. The evolution of $\text{MOI}_{\text{recon}}$ during non-normal amplification. Each curve shows the temporal evolution of the maximally amplifying mode for the labelled norm, with maximum amplification times as indicated.

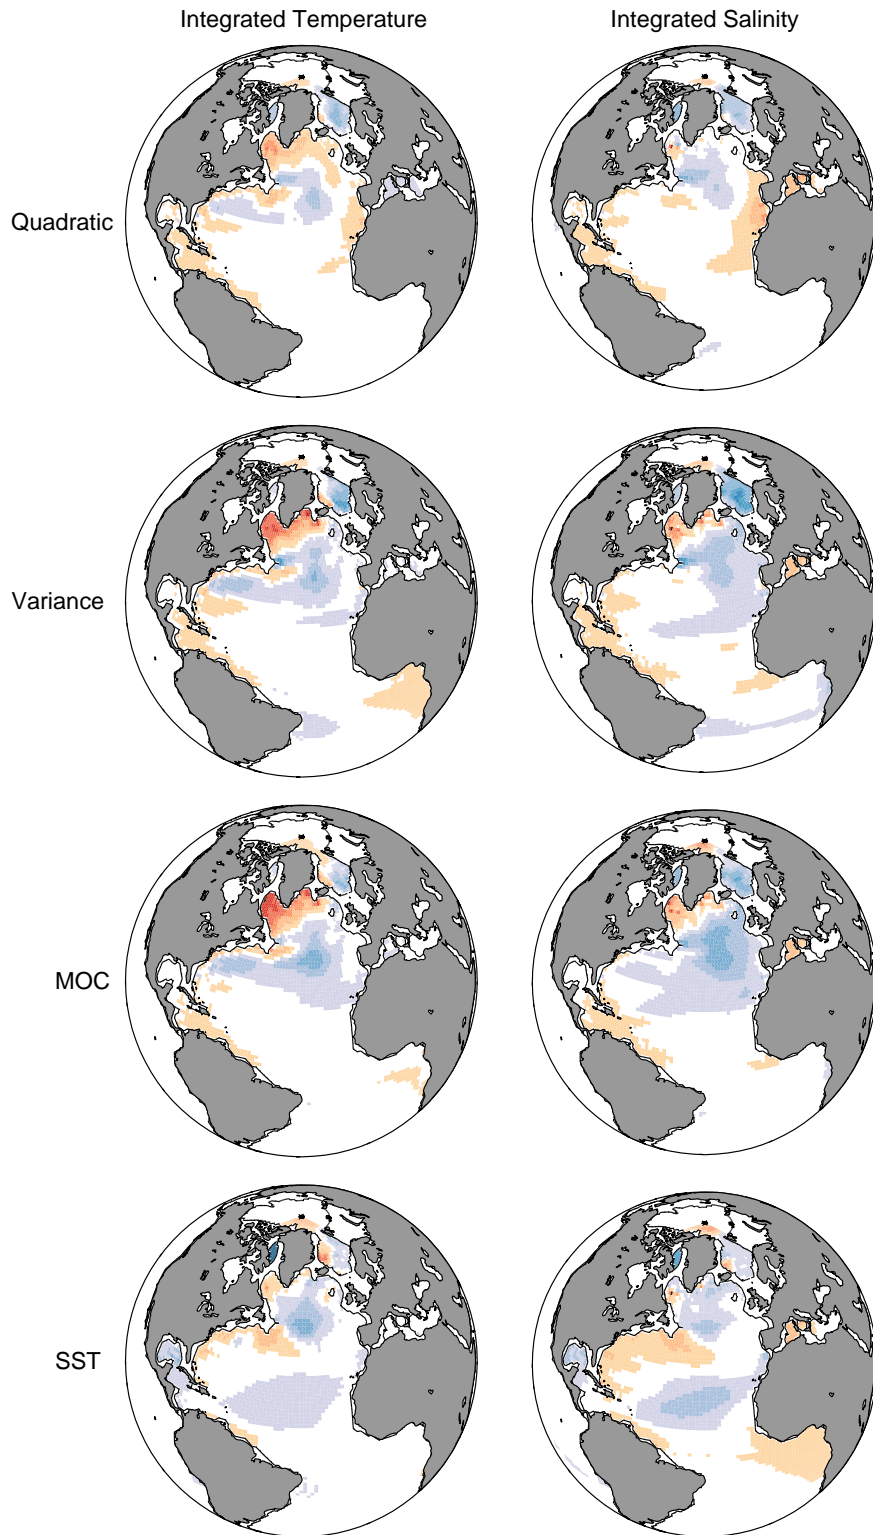


FIG. 11. The initial conditions for maximal amplification over 10 years for each norm, integrated from the surface to a depth of 1800m. Left column: temperature (in K). Right column: salinity (in psu), multiplied by five. The colour scale is the same in every panel and is arbitrary. White regions represent small anomalies of either sign.

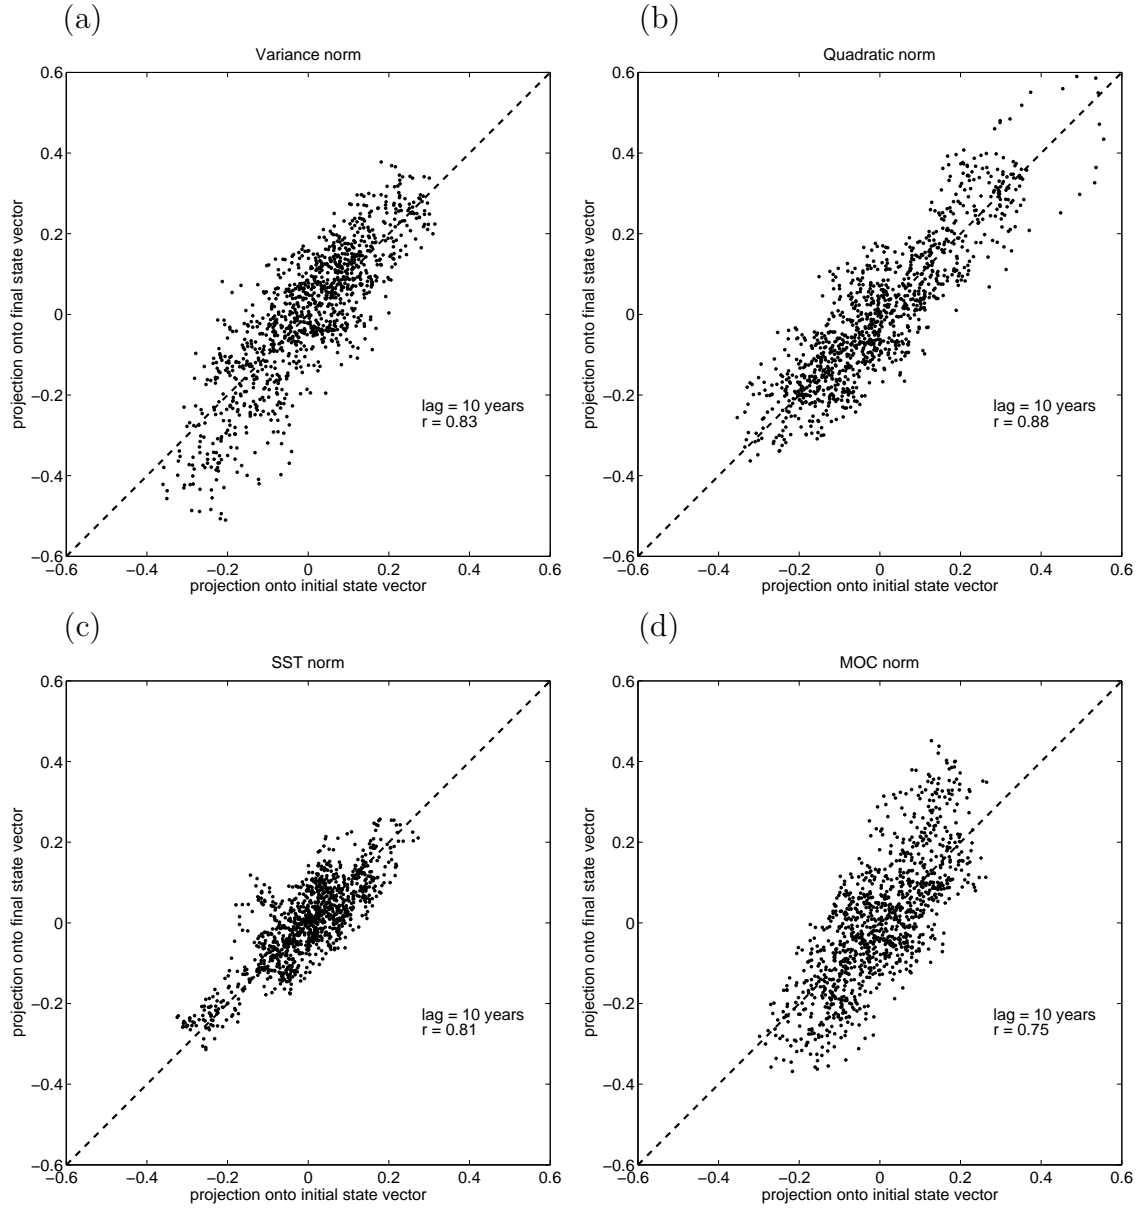


FIG. 12. Scatter plots of the projection of the state vector, $\mathbf{x}(t)$, onto the optimal initial condition, $\mathbf{x}_0(\tau)$, against the projection of $\mathbf{x}(t + \tau)$ onto the maximally amplified state, $\mathbf{P}_\tau \mathbf{x}_0(\tau)$, for all t , and fixed $\tau = 10$ years. The dashed line in each panel has a slope of unity. (a) Using the variance norm. (b) Using the quadratic norm. (c) Using the SST norm. (d) Using the MOC norm.




FIG. 13. Power spectra for the leading five principal components. Solid lines: for 1100 years of the GCM data, x . Dashed lines: ensemble mean of 50 separate segments of 1100 years from a forward model with 95% confidence limits shaded grey.

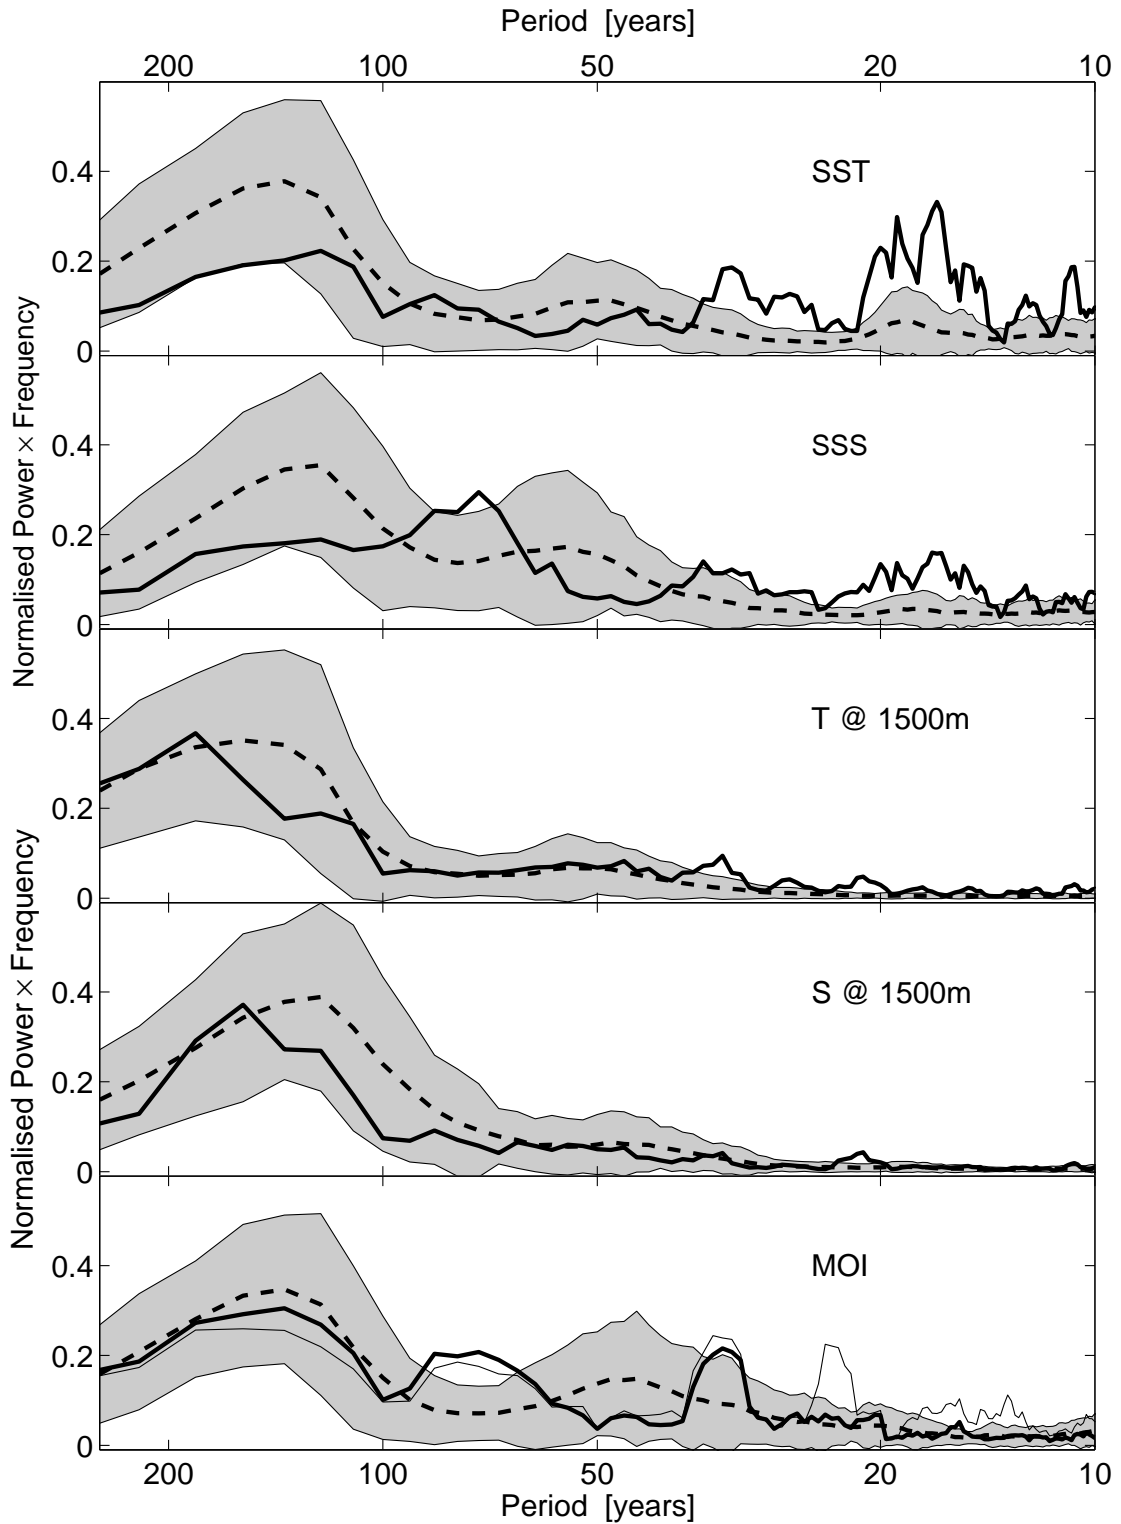


FIG. 14. Power spectra for SST, SSS, temperature and salinity at a depth of 1500m, and the MOI. Solid lines: for 1100 years of the GCM data, \times . Dashed lines: ensemble mean of 50 separate segments of 1100 years from a forward model with 95% confidence limits shaded grey. In the MOI panel the thick solid line represents $\text{MOI}_{\text{recon}}$ and the thin solid line represents MOI_{GCM} .

List of Tables

1	Summary of norms used in the analysis.	40
2	Properties of the Empirical Normal Modes (ENMs) of the 30 EOF system. .	41

TABLE 1. Summary of norms used in the analysis.

Type of norm	Initial norm kernel, \mathbf{L}	Final norm kernel, \mathbf{N}	Maximum amplification	
			time [years]	factor
quadratic variance	\mathbf{I}	\mathbf{I}	35	5.0
	\mathbf{N}_V	\mathbf{N}_V	36	10.2
	\mathbf{N}_{SST}	\mathbf{N}_{SST}	5	32.0
	\mathbf{N}_V	\mathbf{N}_{MOC}	10	1.4

TABLE 2. Properties of the Empirical Normal Modes (ENMs) of the 30 EOF system.

Mode number	Period (years)	Decay time (years)
j	$2\pi/\omega_j$	$-1/\sigma_j$
1	∞	16.3
2	∞	19.6
3/4	4.4	2.8
5/6	9.0	5.6
7/8	10.5	8.2
9/10	14.0	9.9
11/12	19.1	24.1
13/14	27.4	18.4
15/16	30.6	15.3
17/18	50.7	38.7
19/20	74.8	31.2
21/22	75.8	11.3
23/24	134.0	72.2
25/26	157.3	135.1
27/28	185.0	77.3
29/30	414.7	64.2