

Nonnormal Thermohaline Circulation Dynamics in a Coupled Ocean–Atmosphere GCM

ELI TZIPERMAN AND LAURE ZANNA

Department of Earth and Planetary Sciences, and School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts

CECILE PENLAND

NOAA/ESRL/Physical Sciences Division, Boulder, Colorado

(Manuscript received 23 January 2007, in final form 19 July 2007)

ABSTRACT

Using the GFDL coupled atmosphere–ocean general circulation model CM2.1, the transient amplification of thermohaline circulation (THC) anomalies due to its nonnormal dynamics is studied. A reduced space based on empirical orthogonal functions (EOFs) of temperature and salinity anomaly fields in the North Atlantic is constructed. Under the assumption that the dynamics of this reduced space is linear, the propagator of the system is then evaluated and the transient growth of THC anomalies analyzed. Although the linear dynamics are stable, such that any initial perturbation eventually decays, nonnormal effects are found to result in a significant transient growth of temperature, salinity, and THC anomalies. The growth time scale for these anomalies is between 5 and 10 yr, providing an estimate of the predictability time of the North Atlantic THC in this model. There are indications that these results are merely a lower bound on the nonnormality of THC dynamics in the present coupled GCM. This seems to suggest that such nonnormal effects should be seriously considered if the predictability of the THC is to be quantitatively evaluated from models or observations. The methodology presented here may be used to produce initial perturbations to the ocean state that may result in a stricter estimate of ocean and THC predictability than the common procedure of initializing with an identical ocean state and a perturbed atmosphere.

1. Introduction

The present-day North Atlantic Ocean exhibits multidecadal variability in sea surface temperature (SST), possibly partially caused by the large-scale thermohaline circulation (THC) variability (e.g., Bjerknes 1964; Kushnir 1994). Numerous studies have considered the stability and variability of the THC on these time scales using all ranges of model complexity from simple box models to general circulation models (GCMs; e.g., Delworth et al. 1993; Griffies and Tziperman 1995). Studies of the North Atlantic predictability find possible decadal predictability skill of the THC and North Atlantic SST (Collins et al. 2006; Griffies and Bryan 1997).

In previous studies of weather or ENSO predictability, “nonnormal dynamics” and “transient amplification” were found to be crucial (e.g., Farrell 1988, 1989;

Farrell and Ioannou 1996; Penland and Sardeshmukh 1995; Kleeman and Moore 1997). These terms refer to the property of damped (stable) linear systems, which display no variability unless forced externally, to produce strong amplification of appropriate initial perturbations before eventually decaying back to the background state. Such behavior occurs in linear systems in which the operator does not commute with its Hermitian transposed. In other terms, the system $dx/dt = \mathbf{A}x$ is nonnormal if $\mathbf{A}\mathbf{A}^T \neq \mathbf{A}^T\mathbf{A}$, which is the case in practically all systems based on fluid dynamics equations. The nonnormal dynamics of the THC was investigated in only a few recent papers, all using fairly simple models (Lohmann and Schneider 1999; Tziperman and Ioannou 2002; Zanna and Tziperman 2005).

In this work, we study the possible transient amplification of the THC in the Geophysical Fluid Dynamics Laboratory (GFDL) coupled atmosphere–ocean GCM CM2.1 (Delworth 2006; Delworth et al. 2006; Gnanadesikan et al. 2006; Stouffer et al. 2006; Wittenberg et al. 2006; Griffies et al. 2005). The study of nonnormal

Corresponding author address: Eli Tziperman, Harvard University, 24 Oxford St., Cambridge, MA 02138.
E-mail: eli@eps.harvard.edu

dynamics and transient amplification in a coupled GCM introduces nontrivial difficulties. As a first approximation, we construct an EOF-based reduced space based on temperature and salinity fields from a 2000-yr control model run. We then fit a linear model to the reduced space dynamics following the “linear inverse modeling” approach used in the context of ENSO (e.g., Penland and Sardeshmukh 1995; Penland 1996; Penland et al. 2000; Penland and Matrosova 1998, 2001; Blumenthal 1991; Moore and Kleeman 2001). In section 2 we introduce the methodology for constructing the reduced space and the linear fit of the model followed by an analysis of the quality of the fit. We discuss the results of the nonnormal effects in section 3 and finally conclude in section 4.

2. Methodology

To analyze the nonnormal amplification of the THC in the GFDL CM2.1, we first need to obtain a linearized version of the model dynamics. Owing to the difficulty of deriving the explicit linearized equations of the full coupled GCM, we follow the “linear inverse modeling” approach developed in the context of climate research by Penland and coauthors (e.g., Penland and Sardeshmukh 1995; Penland 1996). A reduced space representation of the dynamics is first created by expanding some of the model variables into empirical orthogonal functions and by using the principal component amplitudes as the state vector (Blumenthal 1991). Although principal oscillation pattern analysis (Hasselmann 1988; Von Storch et al. 1988) is the first step of linear inverse modeling, we are not as interested in defining the dominant eigenvectors of the linear feedback matrix as much as we are in finding that matrix itself.

The use of three-dimensional (3D) multivariate EOFs for describing the dynamics and for state space reduction is supported by the recent work of Hawkins and Sutton (2007). Using the Hadley Centre model, they showed that multivariate EOFs provide an effective description for the THC variability and could associate the mechanism for the multidecadal variability to an internal ocean mode, dominated by changes in convection in the Nordic seas, and affected by salinity transport from the Arctic and North Atlantic. An alternative to the 3D EOFs could be the use of dynamic topography, which provides an efficient way of representing the gyre dynamics, as used in the THC predictability study of Griffies and Bryan (1997). From previous studies on nonnormal growth of the THC, we expect the compensation effects of the temperature and salinity to play a significant role in the THC amplifica-

tion. Therefore, the use of dynamic topography, which is based on density alone, is not expected to be optimal for the purposes of this study. However, we still use dynamic topography as a useful diagnostic of the nonnormal dynamics.

Given the reduced space representation, we fit a linear matrix to the evolution of the reduced state vector and analyze its nonnormal properties.

a. Reduced space for the GFDL model

We analyze the nonnormal dynamics of the “1860-control” integration of the coupled GFDL model version CM2.1. Three-dimensional temperature, salinity, and velocity fields are available for this 2000-yr-long control run. The atmospheric model resolution is 2° latitude by 2.5° longitude with 24 vertical levels. The ocean model has a resolution of 1° by 1° poleward of 30° latitude and progressively increased resolution in latitude approaching the equator with $\frac{1}{3}^\circ$ resolution at the equator. The total number of vertical levels is 50, with 22 levels evenly spaced within the first 220 m from the surface. The reader is invited to consult the following papers and references within for a detailed description of the model: Delworth (2006), Delworth et al. (2006), Gnanadesikan et al. (2006), Stouffer et al. (2006), Wittenberg et al. (2006), Griffies et al. (2005). Our main interest in this work is the variability of the North Atlantic THC. Its strength is measured by an index defined as the maximum overturning streamfunction occurring anywhere between 20° and 80°N . The red line in Fig. 1 shows a time series of the THC index deviation from its mean [the mean is 23.9 Sv ($\text{Sv} \equiv 10^6\text{ m}^3\text{ s}^{-1}$) and the standard deviation is 1.5 Sv]. The power spectrum of the THC index demonstrates its variability with two spectral peaks at about 20 and 300 yr and a possibly a weaker peak at about 60 yr (Fig. 2, Delworth and Zhang 2007). We are interested in finding if nonnormal THC ocean dynamics plays a role in this variability.

Our interest being the THC variability on an interannual time scale and longer, we consider annually averaged temperature and salinity fields in the North Atlantic region ($10^\circ\text{--}80^\circ\text{N}$, $20^\circ\text{E}\text{--}100^\circ\text{W}$). The reduced space representation of the GCM dynamics includes only the 3D temperature and salinity fields. To a first-order approximation, we expect the velocity field relevant to the large-scale THC dynamics to be mostly geostrophic (and determined by the temperature and salinity fields) and, therefore, we do not explicitly include the velocities in the state vector.

To reduce the dataset sizes and enable the EOF calculation, we subsampled the temperature and salinity




FIG. 1. Time series of the GCM THC index anomaly in red [where the THC index is defined as the value of the maximum overturning streamfunction (S_v) occurring anywhere between 20° and 80°N , and its anomaly is the deviation from its time mean], and the reconstructed THC index, in green, calculated using the regression coefficients from the temperature and salinity PCs.

fields to a resolution of 3° horizontally (every third model grid point) and 13 levels vertically from 5 to 4950 m (every third vertical level). The temperature and salinity anomalies were obtained by removing their time mean at each grid point. Let the dimensional temperature (similarly for salinity) at a location i and time t be T_{it}^* , and a nondimensional temperature be $\hat{T}_{it} = T_{it}^*/W_i$. The weighting factor W_i may be some relative grid volume $\Delta V_i/V_0 = (dx_i dy_i dz_i)/V_0$, the standard deviation (SD) of the temperature and salinity at a given grid point i , σ_i , or any combination of the two, for example,

$$W_i = \sigma_i/(\Delta V_i/V_0). \quad (1)$$

Note that weighting by the SD tends to make all depth levels more equal, eliminating the dominance of the more energetic surface. Weighting by the grid element volume as in the last equation tends to emphasize the deeper ocean, where grid boxes are larger. Several analyses using different such combinations are discussed below. Moore and Kleeman (2001) have shown that, in order to capture most of the relevant dynamics for the analysis of optimal perturbations, the EOFs should be evaluated from the correlation matrix (or, equivalently, the covariance matrix of the temperature and salinity normalized by their SD) rather than from the dimensional covariance matrix. In addition, before calculating the EOFs, the temperature and salinity fields were smoothed by three passes of a local running horizontal average operator (no temporal smoothing was used beyond the annual averaging). The smoothing results in more easily interpretable 3D structures of the EOFs although there is the danger of misrepresenting small-scale critical dynamics (e.g., the Labrador Sea convection in the current context). It appears that non-smoothed fields result in equally nonnormal operators and comparable transient amplification. Therefore, let the subsampled normalized smoothed temperature and

salinity (\hat{T} , \hat{S}) at all grid points of the model at a given time t_n be represented by a single vector $\mathbf{Y}(t_n)$. The multivariate 3D EOFs of the normalized temperature and salinity are then calculated as the eigenvectors of the covariance matrix $\mathbf{C} = \langle \mathbf{Y}\mathbf{Y}^T \rangle$, where the angled brackets denote time average.

The 3D EOFs generally represent large-scale spatial structures (e.g., Figs. 3, 4), which depend, of course, on the weighting used. Using normalization by the grid volume and by the standard deviation at each grid point tends to emphasize the deep variability, as seen in Fig. 3 for the first EOF. The THC depends on both the deep temperature and salinity as well as on their near-surface signal. The correlation coefficients between the different principal components and the THC index time series (Fig. 5) show that the fifth EOF, for example, plays an important role in setting the amplitude of the THC, and its structure (shown in Fig. 4) is indeed more sur-




FIG. 2. Power spectrum of the THC index time series anomaly obtained from the GFDL coupled GCM control run. The gray shaded region indicates the 95% confidence interval.

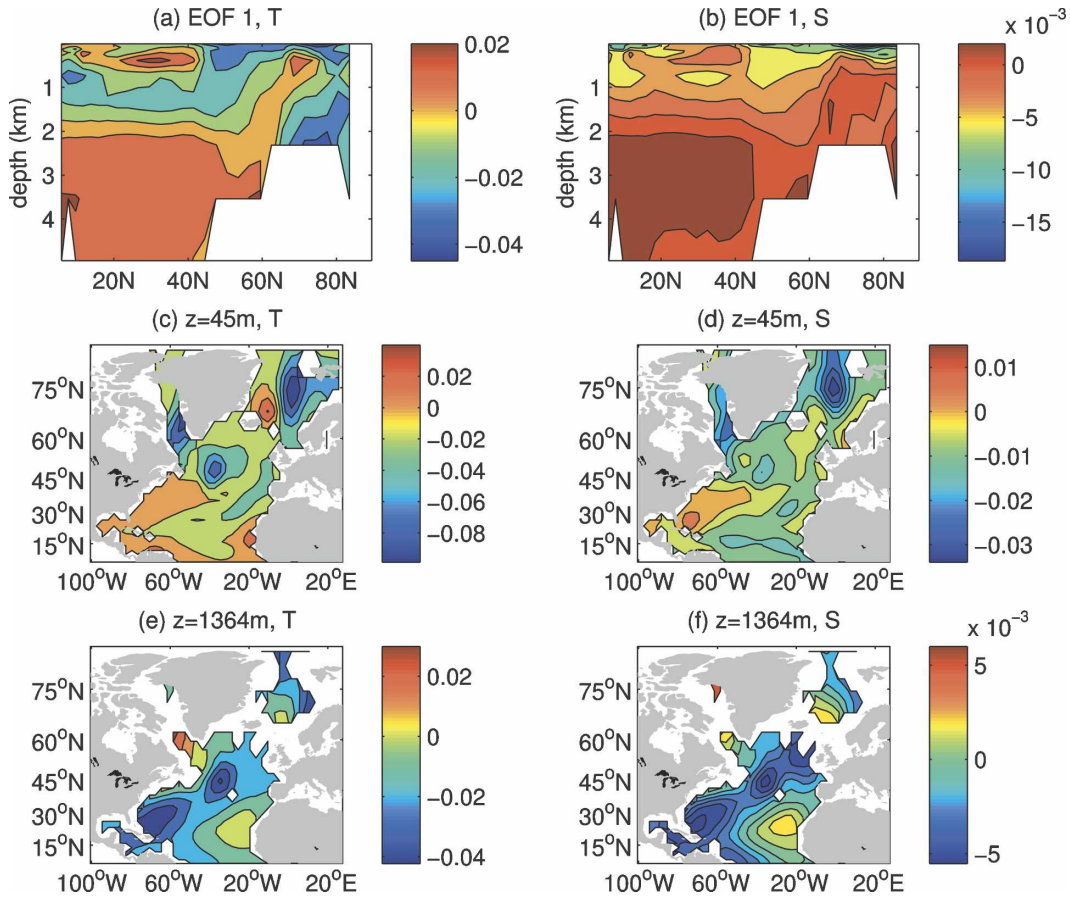


FIG. 3. EOF 1: vertical cross section of the zonally averaged nondimensionalized (a) temperature and (b) salinity as function of latitude and depth; horizontal section at a depth of 45 m of (c) temperature and (d) salinity as a function of longitude and latitude; (e) and (f) as in (c) and (d) but at a depth of 1364 m.

face trapped. EOF 9 (not shown) is also strongly correlated with the THC, yet is not as surface trapped as EOF 5. Changing the weighting factor W_i used to normalize the temperature and salinity results in different EOFs representing this surface-trapped signal and, therefore, in different principal components (PCs) being correlated with the THC index. However, as we will see below, some interesting conclusions regarding the nonnormal dynamics of the THC may be drawn that are not sensitive to the weighting used.

The EOFs are found using the smoothed data fields; however, the principal components are found by projecting the nonsmoothed fields on the EOFs. Effectively, this is a partial smoothing of the data that discards the noisy part of the data that does not project onto the first few EOFs kept for the analysis. By calculating the PC time series for the first N EOFs, we obtained a reduced space representation of the variability, and a new N -dimensional state vector $\mathbf{P}(t_n) = \mathbf{P}_n$. An examination of the PC time series shows a nonneg-

ligible trend in some of the PCs during the first 500 yr. Therefore, we removed the first 500 yr of the data and only use the final 1500 yr in the following analysis. We note that using EOFs as the basis for a reduced space reconstruction is not necessarily the most efficient approach for the purpose of extracting the nonnormal dynamics because EOFs, in general, are optimized for representing the maximum variance rather than the optimal initial perturbations that may lead to large variability (Farrell and Ioannou 2001). We shall briefly return to this issue later.

b. Linear inverse model and propagator

Assuming that the dynamics of reduced space are mostly linear and driven by white noise, we can estimate the propagator $\mathbf{B}(\tau_0)$ of the evolution of the state vector from time n to time $n + \tau_0$,

$$\mathbf{P}_{n+\tau_0} = \mathbf{B}(\tau_0)\mathbf{P}_n, \tag{2}$$

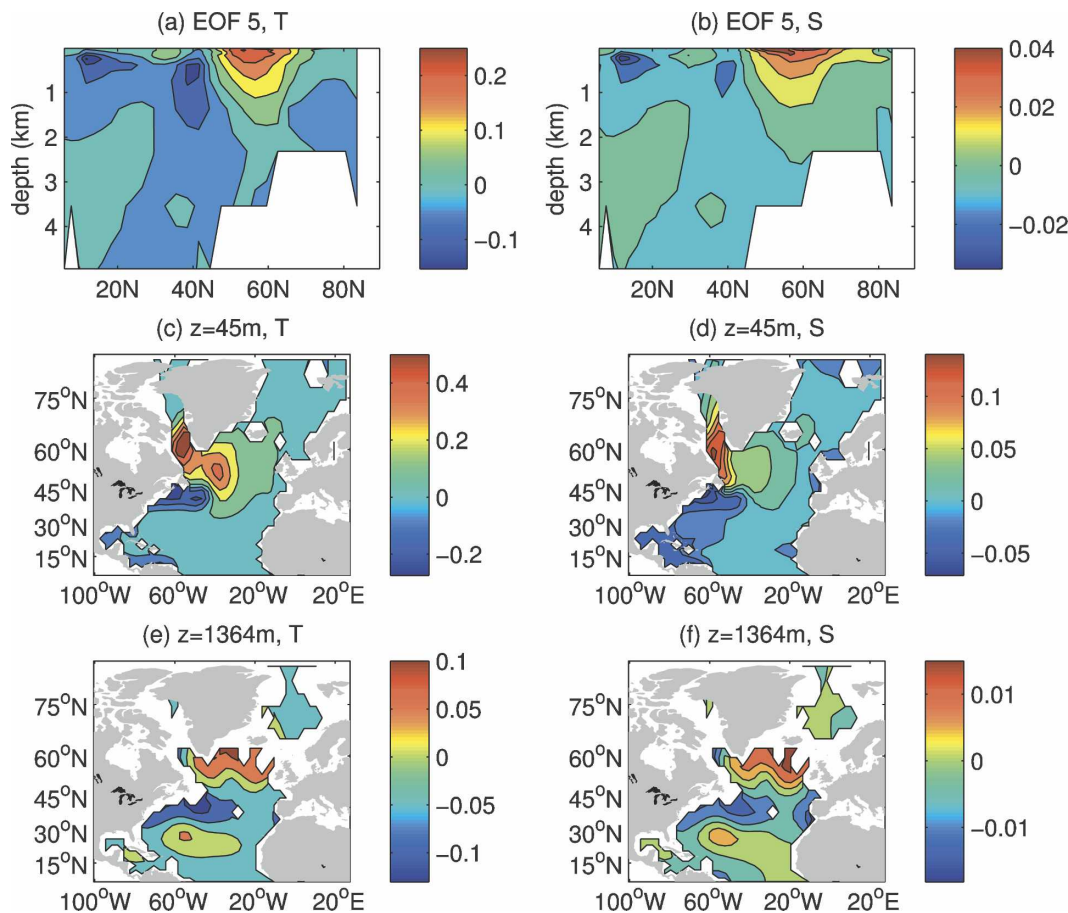


FIG. 4. As in Fig. 3, but for EOF 5.

by minimizing the variance of the residuals given by $\mathbf{r}_{\tau_0} = \mathbf{P}_{n+\tau_0} - \mathbf{B}(\tau_0)\mathbf{P}_n$. The propagator for an arbitrary lead time τ is then calculated as $\mathbf{B}(\tau) = \mathbf{B}(\tau_0)^{\tau/\tau_0}$. The assumption that the evolution of the principal component vector, \mathbf{P}_n , may be described by linear dynamics is anything but obvious. The GCM equations are nonlinear, and it is more than possible that nonlinear effects play a role as well. We nevertheless proceed, keeping in mind this caveat.

The propagator $\mathbf{B}(\tau)$ depends, in principle, on the lag used during the fitting process, τ_0 . We therefore use different values of τ_0 in order to evaluate the propagator in addition to trying different weighting factors W_i and different choices for the number of EOFs N used to construct the reduced phase space. All of the experiments are summarized in Table 1. Unless noted otherwise, the results plotted are from run 1 for which we used 25 EOFs, a lag time of 1 yr, and a weighting by both grid volume and SD.

Once the matrix $\mathbf{B}(\tau)$ is available, it is straightforward to study its nonnormal dynamics, including optimal initial conditions and transient amplification. We

start by defining a norm by which the magnitude of the state vector is evaluated. Let \mathbf{X} be an $N \times N$ positive definite symmetric matrix to serve as the “norm kernel.” The norm of the state vector is then given by

$$\|\mathbf{P}_\tau\|_{\mathbf{X}}^2 \equiv \mathbf{P}_\tau^T \mathbf{X} \mathbf{P}_\tau = [\mathbf{B}(\tau)\mathbf{P}_0]^T \mathbf{X} \mathbf{B}(\tau)\mathbf{P}_0. \quad (3)$$

Looking for the unit norm optimal initial conditions, \mathbf{P}_0 , which maximize the norm of the state at time τ ,




FIG. 5. The correlation coefficients between the THC index anomaly and the PC amplitudes.

TABLE 1. A summary of the different calculations used to linearly approximate the dynamics of the GFDL CM2.1 large-scale North Atlantic temperature and salinity variability (y: yes, n: no). The last line shows the percentage of T, S variance explained by the number of EOFs retained in each experiment.

Parameters	Case												
	1	2	3	4	5	6	7	8	9	10	11	12	13
No. of EOFs	25	25	25	20	30	25	25	25	20	30	25	25	30
Lead time τ_0	1	2	3	1	1	1	2	3	1	1	1	1	2
Weight by SD	y	y	y	y	y	n	n	n	n	n	y	n	y
Weight by dv	y	y	y	y	y	y	y	y	y	y	n	n	y
Var T, S explained	99	99	99	96	99	97	97	97	93	99	95	97	99

$\|\mathbf{P}_\tau\|_{\mathbf{X}}^2$, leads to the following generalized eigenvalue problem for the optimal initial conditions (Farrell 1988; Tziperman and Ioannou 2002):

$$\mathbf{B}(\tau)^T \mathbf{X} \mathbf{B}(\tau) \mathbf{P}_0 = \lambda \mathbf{X} \mathbf{P}_0. \quad (4)$$

We will consider two different norms: First, the energy norm, where \mathbf{X} is the identity matrix, and then a norm kernel that measures the amplitude of the THC. To construct the THC norm kernel we first find the regression coefficients \mathbf{R}_{THC} between the PCs and the THC index by minimizing the sum (over time) of the square of the residuals

$$r(t_n) = \text{THC}(t_n) - \mathbf{R}_{\text{THC}}^T \mathbf{PC}(t_n). \quad (5)$$

The correlation between the THC index and the THC reconstructed using the regression from the PCs is approximately 0.88, indicating that the PCs are an appropriate state vector for studying the THC dynamics, and the two are shown in Fig. 1. The good correlation between the THC linearly reconstructed from the PCs and the original THC index indicates that the THC is linearly related to the principal components and hence to the temperature and salinity fields. While this is encouraging as far as our linear framework is concerned, it still does not assure us that the PC dynamics are linear as well and may be described by a linear model such as (2).

Given the regression coefficients \mathbf{R}_{THC} , we define the THC norm kernel matrix as

$$\mathbf{X}_{\text{THC}} = \mathbf{R}_{\text{THC}} \mathbf{R}_{\text{THC}}^T. \quad (6)$$

The correlation coefficient for each PC is shown in Fig. 5. When using this norm kernel in (4), we attempt to find the initial conditions in the reduced EOF space that maximize $\text{THC}^2(t = \tau)$ subject to the condition that $\text{THC}^2(t = 0) = 1$.

Before proceeding with the analysis of the fit, we should mention that this norm kernel is singular, being defined using a single vector, and it therefore requires regularization and careful interpretation (Tziperman

and Ioannou 2002; Zanna and Tziperman 2005). The regularization is done by adding a regular norm kernel multiplied by a small number. That is, \mathbf{X}_{THC} is replaced by $\mathbf{X}_{\text{THC}} + \varepsilon \mathbf{I}$, for example, if the identity matrix is used. The largest eigenvalue of the nonregularized problem is infinite, and the corresponding initial THC is zero. As ε approaches zero, the eigenvalues of the regularized problem approach infinity, but the eigenvectors (optimal initial conditions) do not change as long as ε is not too small to cause numerical problems. However, the eigenvectors are sensitive to the form of the added norm kernel. If we add a matrix other than the identity matrix, the eigenvectors will be different. This is not surprising, given that optimal initial conditions are, in general, sensitive to the norm kernel used. The regularizing norm kernel constrains, in a sense, the null space of the THC singular norm kernel, and thus the resulting eigenvectors depend on its choice. Overall, it seems that the resulting eigenvectors have a large scale, seemingly physical structure, and they result in growth of the THC as expected. We feel that the approach used here of using the regular energy norm kernel on the one hand, and the singular but regularized THC norm kernel on the other, provides a useful information on the nonnormal THC dynamics. An alternative approach to dealing with this singularity, based on using a different norm (L_∞ rather than L_2) is explored in Zanna and Tziperman (2008).

c. Quality of fit

Before studying the nonnormal dynamics of the linear propagator, we need to verify that our fit does not violate any of our assumptions and that the results are not overly sensitive to the different parameters used. In general, we found the linear fit (2) to the GFDL CM2.1 to be relatively well behaved and to result in a physically consistent propagator for $\tau \leq 2$ yr (for $N = 20, 25$, or 30 EOFs).

More specifically, a physical linear propagator of a stable linear system needs to have eigenvalues with

magnitude smaller than one and with a positive real part. Using $N = 25$, the smallest real part of the eigenvalues of \mathbf{B} (τ_0) for $\tau_0 = 1$ yr is 0.77. The lag time $\tau_0 = 2$ yr still provides physical results, with smaller eigenvalues reflecting the further decay of the modes in the additional year. With $\tau_0 = 3$ yr, two of the eigenvalues turn out to have very small negative real parts, which is unacceptable and indicates that the fit has failed because the short decay time of some of the modes cannot be resolved using a 3-yr lag time.

The quality of the fit can be further examined by comparing the predictability using the fitted model for different values of the lag time τ to the expected predictability of a perfect linear model driven by white noise, and to the predictability of an autoregressive model $O(1)$ (AR1). The latter is obtained by fitting a diagonal \mathbf{B} to (2). In addition, the dynamics should be independent of the value of τ_0 used to obtain the propagator (“tau test,” Penland 1996).

Consider now the prediction error covariance matrix for a perfect model. Let $\hat{\mathbf{P}}_n$ be the actual state derived from the GCM EOFs. The residuals of the prediction are given by

$$\mathbf{r}_\tau = \hat{\mathbf{P}}_{n+\tau} - \mathbf{B}(\tau)\hat{\mathbf{P}}_n, \quad (7)$$

and the prediction error covariance can be shown (Penland 1989) to be

$$\begin{aligned} \mathbf{C} &= \langle \mathbf{r}_n \mathbf{r}_n^T \rangle = \langle (\hat{\mathbf{P}}_{n+\tau} - \mathbf{B}\hat{\mathbf{P}}_n)(\hat{\mathbf{P}}_{n+\tau}^T - \hat{\mathbf{P}}_n^T \mathbf{B}^T) \rangle \\ &= \mathbf{C}_P - \mathbf{B}\langle \hat{\mathbf{P}}_n \hat{\mathbf{P}}_{n+\tau}^T \rangle - \langle \hat{\mathbf{P}}_{n+\tau} \hat{\mathbf{P}}_n^T \rangle \mathbf{B}^T + \mathbf{B}\mathbf{C}_P\mathbf{B}^T \\ &= \mathbf{C}_P - \mathbf{B}\langle \hat{\mathbf{P}}_n \hat{\mathbf{P}}_n^T \rangle \mathbf{B}^T - \langle \hat{\mathbf{P}}_{n+\tau} \hat{\mathbf{P}}_{n+\tau}^T \rangle \mathbf{B}^T + \mathbf{B}\mathbf{C}_P\mathbf{B}^T \\ &= \mathbf{C}_P - \mathbf{B}\mathbf{C}_P\mathbf{B}^T, \end{aligned} \quad (8)$$

where \mathbf{C}_P is the covariance matrix of the actual state vector from the GCM and is given by

$$\mathbf{C}_P = \langle \hat{\mathbf{P}}_n \hat{\mathbf{P}}_n^T \rangle = \langle \hat{\mathbf{P}}_{n+\tau} \hat{\mathbf{P}}_{n+\tau}^T \rangle. \quad (9)$$

This “perfect model” prediction error covariance \mathbf{C} given in (8) is compared to the actual prediction error covariance obtained by predicting the PC state vector using the fitted linear propagator calculated from (2) as well as to the prediction error covariance obtained from the AR1 model. The measure to be compared may be the trace of the error covariance matrix (Fig. 6b), or the individual diagonal elements (Figs. 6c–i). The prediction skill of our fitted linear model is generally better than that of the AR1 model (which is a stricter test than the commonly used comparison to persistence), although not as good as a perfect linear model is expected to be. We conclude that the fitted linear model is a useful approximate representation of the large-scale T ,

S dynamics in the GCM. This does not rule out a role for nonlinear processes, but indicates that the linear framework may, at least, be expected to be useful.

A similar derivation applies to the covariance of the THC prediction residuals. Define the residuals as

$$r_\tau^{\text{THC}} = \text{THC}_{n+\tau} - \mathbf{R}_{\text{THC}}^T \hat{\mathbf{B}} \hat{\mathbf{P}}_n, \quad (10)$$

their covariance, assuming a perfect linear model as well as that the THC is perfectly predicted from the PC amplitudes, $\text{THC}_n = \mathbf{R}_{\text{THC}}^T \hat{\mathbf{P}}_n$, is then

$$C^{\text{THC}} = \langle (r_\tau^{\text{THC}})^2 \rangle = \langle \text{THC}_n^2 \rangle - \mathbf{R}_{\text{THC}}^T \mathbf{B} \mathbf{C}_P \mathbf{B}^T \mathbf{R}_{\text{THC}}. \quad (11)$$

Figure 6a shows the prediction error covariance for the THC based on the fitted linear model, for an AR1 model, and for a perfect linear model (11). The fitted linear model does better than the AR1 model for the first 10 to perhaps even 20 yr, beyond which there is not much of a skill. If we compare this result to Fig. 6b, which shows the trace of the prediction error covariance for the PCs, we see that the fitted linear model for the PC components does better than the AR1 model for a longer time. We note that the assumption of a linear Markov model may not be a completely accurate representation of the coupled model North Atlantic dynamics and can lead to some of the prediction errors observed. Some PCs (e.g., PC1, Fig. 6c and upper panel of Fig. 7) are better predicted than others (e.g., PC5, Fig. 6g and third panel of Fig. 7). Note that the THC is strongly correlated with PC5, which is not well predicted beyond 10 to 20 yr, like the THC itself. This may be a good place to note that PC1 and PC2 seem significantly correlated at a nonzero lag, which may be a reflection of the oscillation mechanism and could be further explored, as done, for example, in Fig. 5 of Hawkins and Sutton (2007).

d. Sensitivity

It is convenient to examine the sensitivity of our results to the various assumptions made by examining, for each of the experiments summarized in Table 1, the maximum amplification, λ , of $\|\mathbf{P}(\tau)\|_{\mathbf{X}}^2$ as function of τ calculated using (4). These maximum amplification curves are shown in Figs. 8 and 9 for the energy norm and the THC norm respectively.

Figure 9 shows that the maximum THC amplification occurs between 5 and 7 yr. This result seems robust to an increase in the number of EOFs N (e.g., cf. runs 4, 1, and 5 or run 1 with run 13 in Fig. 9). Similar results are obtained if we fix both the number of EOFs and the weighting and vary the lag time τ_0 (e.g., cf. runs 3, 2, and

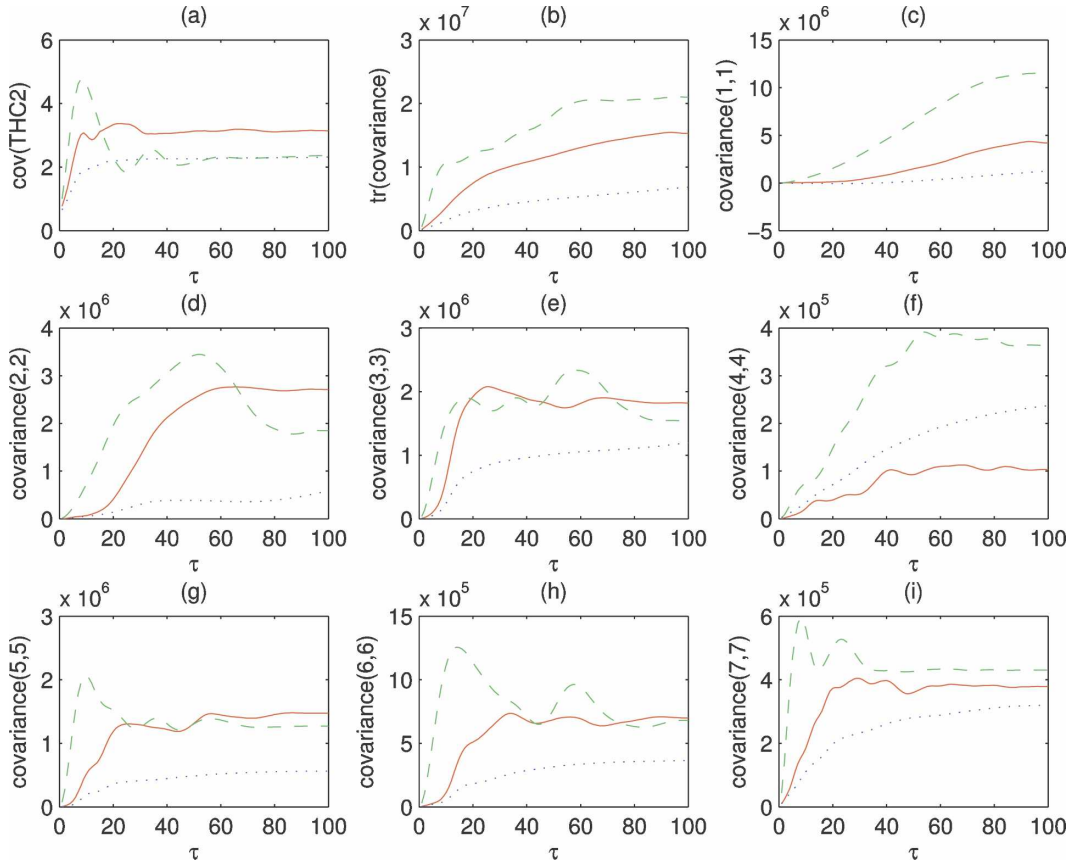


FIG. 6. Quality-of-fit test for run 1 showing the error covariance matrices as a function of forecast lead time for a perfect linear model [Eq. (8), dotted], for an AR1 model (dashed) and for the fitted linear model (solid). (a) The THC error covariance; (b) the trace of the error covariance matrix; (c)–(i) first few diagonal elements of the error covariance matrix, corresponding to PCs 1 to 7.

1 in Fig. 9). Some dependence of the time of maximum amplification on the weighting W_i is also seen on this plot (see in particular run 12, which uses no weighting by SD or grid volume).

The maximum amplification curves using the energy norm, representing the predictability of the PC components rather than of the THC, are shown in Fig. 8. These curves show again a rapid increase in maximum amplification during the first 10 yr or so. Beyond this point, there is a wide spread of the curves, depending mostly on the weighting strategy used.

The smallest amplification, which occurs about 7 yr after the initial conditions, occurs when no weighting is used or when the weighting is based only on the grid volume or only on the SD of the temperature and salinity variability, but not on both. Using a weighting factor that depends on both grid element volume and SD results in the largest transient amplification. The importance of weighting by the standard deviation was also pointed out in the context of ENSO Moore and Kleeman (2001). In addition, the results are sensitive to

the number of EOFs used for the reduced state reconstruction. Using 30 EOFs results in the largest amplification, while 20 and 25 EOFs results in significantly less amplification. Unfortunately, we cannot use a larger number of EOFs due to the short time series available. With 1500 yearly snapshots, we can produce that many equations for the coefficients of the linear model. Using 30 EOFs results in a linear propagator matrix \mathbf{B} whose dimensions are 30×30 , implying 900 elements that need to be solved using these 1500 equations. The resulting conditioning number for the coefficients of \mathbf{B} is reasonable (less than 100), but a further increase in the number of EOFs will quickly result in more unknowns than equations and therefore in an ill-posed problem for the coefficients of the propagator. It may be possible to use the monthly model output (after removing monthly climatology) rather than the yearly averaged fields. But, while this will increase the number of equations, it may not add independent information for constraining the propagator on decadal time scales, in which we are interested. In addition, using monthly




FIG. 7. The 1500-yr time series of different nondimensional principal components from the GCM, together with short segments trying to predict each of them using our linearized model: (top) PC 1, (second) PC 2, and (third) PC 5; the segments are plotted every 10 yr and extend for 150 yr each. (bottom) A 500-yr portion of the THC time series (Sv) with segments plotted every 10 yr and extend for 10 yr each.

data may require making the propagator monthly as well to represent the additional seasonally dependent anomaly dynamics, and for these reasons we have chosen not to explore this direction.

There is also some sensitivity to the lag time τ_0 used in (2) to calculate the fitted propagator $\mathbf{B}(\tau)$ (e.g., cf. curves corresponding to runs 1–3 in Fig. 8). For large enough lead time τ in Fig. 8 the maximum amplification decays to zero, as expected. The eventual decay time scale of the temperature and salinity anomalies and then of the THC is influenced by the decay time scale of the slow decaying modes found to be on the order of 1200 yr for the slowest mode). When using a lag time of $\tau_0 = 1$ or 2 yr to evaluate the propagator [Eq. (2)], it seems likely that such a long time scale for the slow decaying eigenmodes would not be well resolved, and we therefore consider this part of the fit less reliable. As

a result of these sensitivities, in particular to the number of EOFs used, the amplification factors found using the energy norm should be viewed as a lower bound on the expected nonnormal amplification rather than a precise estimate. Also, as mentioned above and discussed in Farrell and Ioannou (2001), using the EOFs for the reduced space reconstruction is not optimal and may again result in an underestimate of the amplification, although some of this problem may have been corrected by the use of the normalized temperature and salinity as discussed above, following Moore and Kleeman (2001).

Some of the variations between the different sensitivity runs may be rationalized as follows: The largest growth in the energy norm is obtained for the largest number of retained EOFs and when the EOFs are normalized by the standard deviation. This underscores the




FIG. 8. Maximum amplification λ , as calculated from Eq. (4), as function of τ for all runs listed in Table 1 using the energy norm.




FIG. 9. As in Fig. 8, but for all runs using the THC norm kernel.

importance of the stochastic forcing by low-variance, unresolved nonlinearities that are weighted up by the normalization to the evolution of the THC. Also, the most variation among the runs with $\tau_0 = 1, 2,$ and 3 yr is found when the EOFs are normalized by the standard deviation. This is evidence that these unresolved nonlinearities are the reason for failing the τ test. In contrast, the maximum amplification curves for runs 6–10 (dashed curves in Fig. 8), for which no normalization by the SD is done and which therefore does not include the low-variance unresolved nonlinearities, pass the τ test (i.e., similarity of results for different values of τ_0) reasonably well in the energy norm. Thus, it seems that the larger resolved scales behave more linearly than do the small unresolved scales. When the PCs are normalized, the small scales dominate the growth, and this growth decreases with τ_0 (cf. runs 5 and 13 in Figs. 8 and 9). The fewer EOFs are retained, the more signal is treated as unpredictable stochasticity. The linear model parameterizes nonlinearities as a linear part plus a stochastic part. When the fast nonlinear signals are included in the mix of EOFs and τ_0 is small enough to resolve their dynamics, they show up as added growth, essentially by putting a best-fit linear model to the nonlinear part.

One important sensitivity, which is not explored here, has to do with the model domain, which was somewhat arbitrarily chosen as small as possible here, to minimize the problem size and enable the calculation of the EOFs. The results may be affected by the choice of the domain. For example, if the domain was to include the Arctic Ocean, the EOFs and the nonnormal growth effects may have been different, given that several GCMs have found significant variability in the Arc-

tic possibly affecting the THC (e.g., Jungclaus et al. 2005).

We conclude that the fitted linear dynamics and deduced transient amplification depend on some of the choices we make during the fit. The interpretation below would need to carefully consider these sensitivities to some of the fit parameters before proceeding with a physical interpretation and conclusions. We will see that some of the sensitivities indicate possible limits to our analysis due to the available model output and other factors. However, other sensitivities, especially to the weighting used, reflect actual physical effects that should be expected, rather than a problem with our fit.

3. Results

a. Maximum amplification curves

As seen above, the maximum amplification using the energy norm occurs after 7 to 50 yr, depending on the assumptions used during the construction of the linear fit (Fig. 8). The nonnormal amplification of initial temperature and salinity anomalies under the energy norm is quite impressive. Figure 8 shows an amplification of initial anomalies in the PC space by a factor of 5–10 within 8–40 yr. All curves show a rapid amplification during the first 10 yr followed by a wide spread of behavior of the amplification curves beyond this time scale. The amplification curves from all experiments using the THC norm kernel (Fig. 9) have a more uniform overall appearance. All such amplification curves peak after about 5–8 yr and then decay fairly rapidly.

The rapid amplification during the first decade seen in both the THC norm and the energy norm indicates that temperature, salinity, and THC anomalies are rapidly amplified during this initial period. From a predict-

ability point of view, this means that errors in the initial conditions, which partially project on the most rapidly growing modes, are expected to rapidly grow as well during these first 10 yr, implying that much of the predictability skill is expected to be lost over this time interval. Even though the amplification factors for the THC and energy norm kernel (Fig. 8) are more or less sensitive to the different parameters used, it is important to note that all curves show a rapid growth phase during the first 10 yr. It seems, therefore, that the rapid growth of anomalies and expected loss of predictability of the THC, temperature, and salinity on a time scale of about 8 yr is a robust result of this study.

The actual amplification values for the THC norm kernel are large (amplification factor of a few hundreds). In the simpler model studies of nonnormal THC amplification (Tziperman and Ioannou 2002; Zanna and Tziperman 2005), the amplification of THC was found to be infinite as the initial THC anomaly induced by the optimal initial conditions for the temperature and salinity typically vanished. This is the result of a cancellation of the initial temperature and salinity anomalies and occurs mathematically due to the singularity of the THC norm kernel (Tziperman and Ioannou 2002). The large THC amplification seen in this study is due to the same effect.

b. Predictability

The fitted linear model is able to successfully predict the THC up to about 7 yr. Predictability of both temperature and salinity and of the THC seems limited by the initial 5–7-yr-long period of fast amplification seen in Figs. 8 and 9. These figures show that an optimal initial perturbation will be amplified by a significant factor within a period of 5–7 yr. Thus, errors in the initial state that project on the optimal initial perturbations may result in a deviation of the predicted fields from their true evolution within about 5–7 yr. While there is some predictability on a short time scale of a few years, the linear model is unable to predict the longer-term oscillations corresponding to the spectral peak of the THC index at 200–300 yr. Interestingly, the linear model does well predicting the first and even the second EOFs (Figs. 6c,d and 7a,b), which have longer time scales. But this does not translate into a corresponding skill in predicting the slow component of the THC oscillations (as shown by the comparison to the AR1 model in Fig. 6a). The reason for this short predictability time may very well be the nonnormal amplification of initial anomalies to the temperature, salinity, and hence THC. We conclude that the predictability of THC anomalies seems to be limited to the order of 5–7 yr.

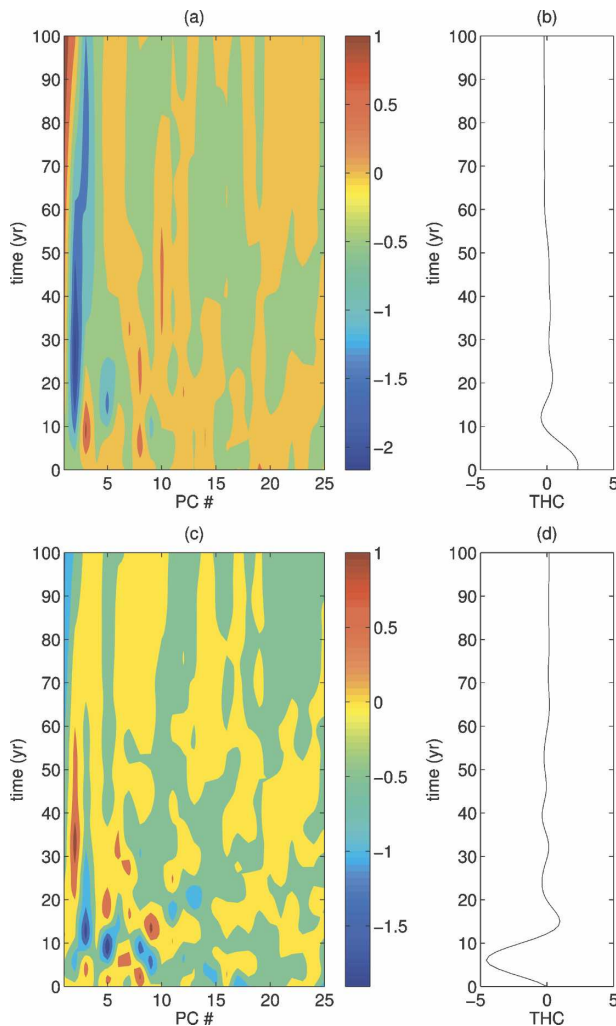


FIG. 10. Evolution in PC space of the fitted linear model from optimal initial conditions and the corresponding THC anomaly for (a), (b) the energy norm and (c), (d) the THC norm. Nondimensional units.

c. Optimal initial conditions and evolution to maximally amplified state in reduced space

The optimal initial conditions and their evolution in PC space for the energy norm and the THC norm are plotted in Figs. 10a and 10c. The corresponding THC anomaly is shown in Figs. 10b and 10d. Under the energy norm, PC 2 is significantly amplified, as are some other PC amplitudes. The THC norm, however, leads to the growth of PCs 2, 5, 8, and 9, some of which are strongly correlated with the THC index as seen in Fig. 5. As expected, the initial THC anomaly vanished under the THC norm kernel (Fig. 10d), but not under the energy norm (Fig. 10b).

Under the THC norm kernel, the optimal initial conditions lead to a decaying oscillation of the THC (Fig.

10d). The maximum growth seems to occur at the first maximum of the oscillation. One wonders if the non-normal analysis simply calculates initial conditions that excite this oscillation and that nonnormal effects do not actually play a role here. However, examining the projection of the optimal initial conditions on the eigenmodes of the propagator \mathbf{B} we find that many more than just a single oscillatory mode are involved in the optimal initial conditions, indicating that nonnormal effects do play a role.

The optimal initial conditions under the energy and THC norms project differently on the PC state vector and therefore involve different spatial structures of the temperature and salinity initial conditions as dictated by the corresponding EOFs. The evolution from the optimal initial conditions seen in the Fig. 10 are very different again, reflecting a possibly different amplification dynamics. That is, different physical processes are responsible for the growth in these different cases. This is consistent with the studies using simpler models, where the physical processes may be more clearly elucidated (e.g., Zanna and Tziperman 2005). In these simpler studies one can attribute the growth to the advection of temperature or salinity anomalies by the mean circulation or the advection of the mean temperature or salinity by the anomalous circulation, etc. Our current approach, constrained by the GCM complexity to using a reduced state vector, does not allow us to point to such specific physical processes.

d. Optimal initial conditions and maximally amplified states in physical space

The amplification of the optimal initial conditions when using the energy norm involves some significant redistribution of temperature and salinity anomalies in space, as can be seen by examining the evolution from the optimal initial conditions to the state at time of maximum amplification (Fig. 11). The plotted density anomalies ρ' are calculated from the temperature and salinity anomalies T' , S' corresponding to the optimal initial conditions or amplified state using the long-term average temperature and salinity in the GCM solution, T_0 , S_0 , using the nonlinear equation of state such that $\rho' = \rho(T_0 + T', S_0 + S') - \rho(T_0, S_0)$. The initial state is very much surface trapped, while the evolved state, already at year 6 or 8, involves temperature and salinity anomalies that are spread over a wider depth range. This provides some physical insight regarding the mechanism of amplification and decay of the optimal perturbations. Note that when weighting by neither grid volume nor SD (runs 6–12, Fig. 8), the perturbations decay fast, after about 20 yr. With weighting by both grid volume and SD [Eq. (1), runs 1–5 and 13, Fig.

8], the amplification takes a longer time and the decay is much slower. The reason for this is rather simple. When the weighting includes these two factors, it tends to emphasize the deep ocean. This means that the optimization is required to produce initial conditions that lead to a growth of temperature and salinity anomalies in the deep ocean. Such anomalies take a bit longer to develop, hence the slower growth time to maximum amplification. The deep anomalies take, of course, much longer time to dissipate than surface anomalies, hence the much longer decay time seen for the runs with weighting by both SD and grid volume. The mathematical machinery of the transient amplification clearly results in physically sensible solutions that are consistent with our intuition of the ocean dynamics, providing another indication that the results are likely to be robust in spite of our many simplifying assumptions.

It is interesting in this context that the overall structure of the amplification of the THC anomalies as a function of time is not as sensitive to the weighting and tends to decay after 10–20 yr for all weighting options used. This suggests, perhaps, that the THC mostly depends on the near-surface (upper 1 km) temperature and salinity anomalies and is therefore not affected by the long-lived deep temperature and salinity anomalies discussed above.

In the horizontal plane, the optimal initial conditions show a warm and salty anomaly in the Labrador Sea that results in a dense anomaly there. This anomaly possibly results from some role played by the Labrador Sea deep convection and its known effect on North Atlantic Deep Water and the amplitude of the THC. In the amplified state, after 27 yr (Fig. 11) the anomaly in the Labrador sea vanishes and appears at 45°N in the west Atlantic. It is also interesting to see the strong density anomaly that develops north of 60°N at all depth levels, driven mostly by the salinity anomaly. Overall the structure of the optimal initial conditions is very different from that of the maximum amplification state, indicating some complex dynamics that does not involve growth of anomalies in place but rather some advection and nonlocalized growth.

Although the maximum amplification of the energy norm occurs after 27 yr, the temperature and salinity evolutions show a clear development of strong perturbations already by year 6, consistent again with the predictability time scale for the THC. It seems that the time scale for nonnormal growth in this model is quite robustly 6–8 yr regardless of the norm kernel used.

The temperature anomaly for the optimal initial conditions and the maximally amplified state after 6 yr under the THC norm (Fig. 12) are both states charac-

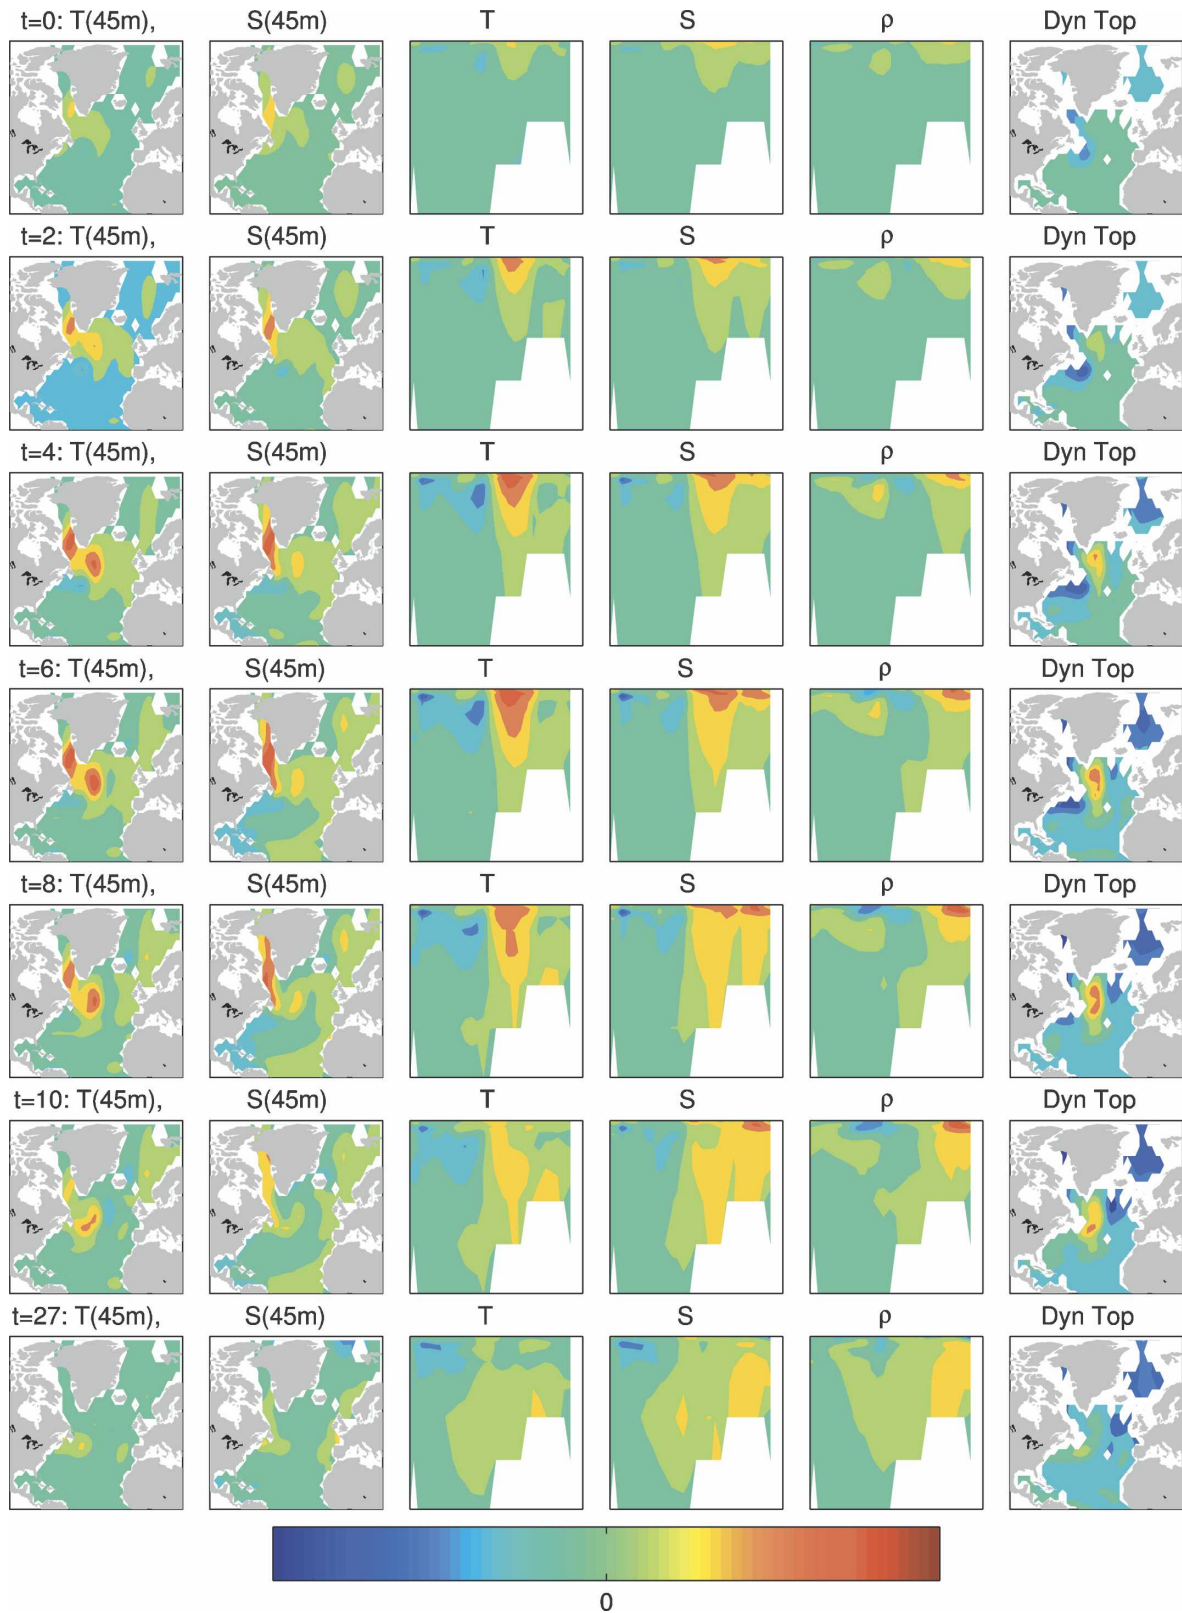


FIG. 11. Evolution from optimal initial conditions using the energy norm kernel leading to a maximum amplification after 27 yr. Each row corresponds to a different time in years as noted in the panel titles. From left to right in each row: temperature at 45-m depth; salinity at 45 m; vertical cross sections of the zonally averaged temperature, salinity, and density; and surface dynamic topography relative to 1364 m. Nondimensional units; domain boundaries as in Fig. 3.

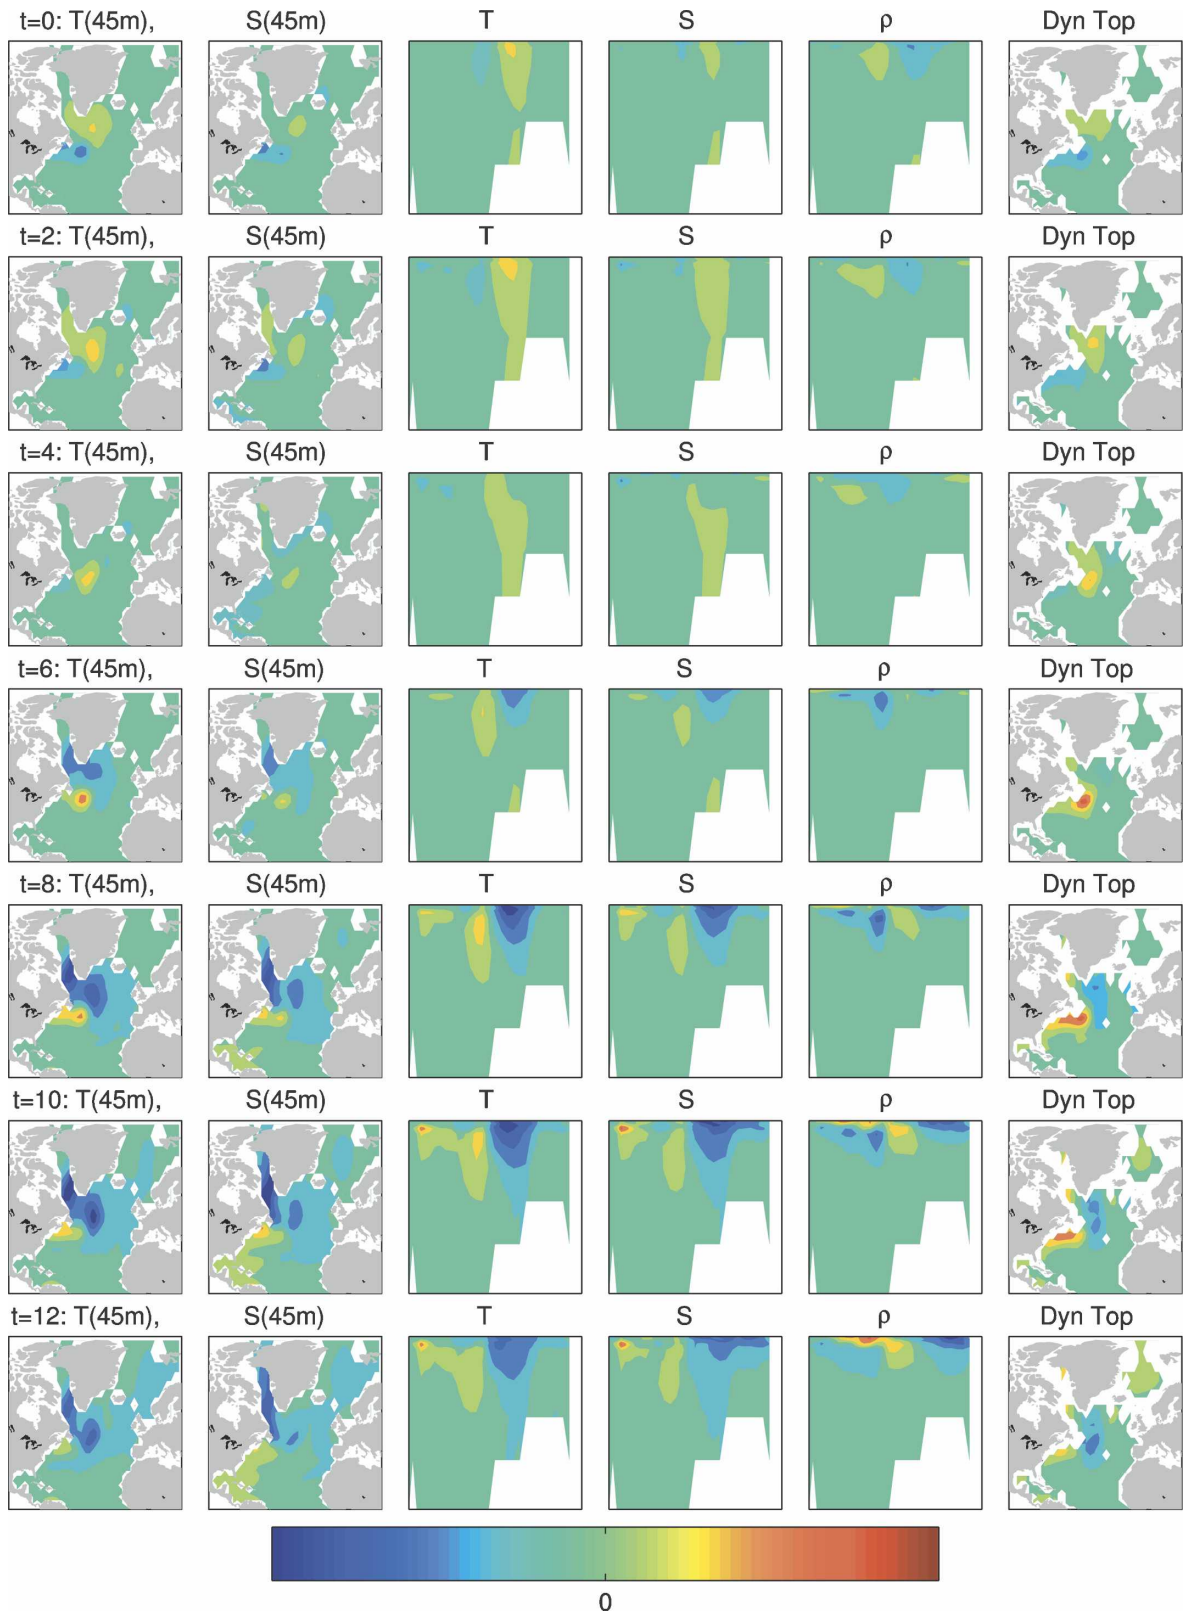


FIG. 12. Evolution from optimal initial conditions, as in Fig. 11, except using the THC norm kernel leading to a maximum THC amplification at 6 yr.

terized by a dipole structure between 40° and 60°N, although the sign of the anomalies in the dipole reverses during the time evolution. The salinity shows a near-surface signal in the optimal initial conditions that develops into a significant anomaly in the amplified state, although it seems that the density anomaly in the amplified state, ultimately responsible for determining the velocity field and therefore the THC amplitude, is dominated by the temperature anomaly rather than the salinity.

The physics of THC oscillations in some GCMs has been shown to involve gyre changes as well (e.g., Delworth et al. 1993), although these gyre changes may not be essential to the very existence of these oscillations (Griffies and Tziperman 1995). Understanding the THC in the currently analyzed model to the degree analyzed in the above papers or by Hawkins and Sutton (2007) is not within the scope of this paper, but we do want to gain some insight into the growth mechanism. Following Griffies and Bryan (1997) we use dynamic topography in order to identify some aspects of the dynamics (rightmost panels in Figs. 11 and 12). The optimal initial conditions under both norms show a weak dynamic topography signal off Newfoundland. This later develops into a rather different yet strong signal for the two different norms (cf. year 10 in Figs. 11 and 12). The gyre signal decays before the maximum amplification time of 27 yr in the energy norm, but seems to be getting stronger after the maximum amplification time of 6 yr in the THC norm. These observations raise the possibility that the gyre is involved in the growth as well as decay of the perturbations. The different evolution in the horizontal plane under the two norms again suggests different physical processes involved. Further understanding of the role of the horizontal gyres would need to await a GCM-based analysis rather than the reduced space description used here.

e. Discussion

One of the more interesting insights here has to do with the difference from previous studies using simpler models. As explained above, a singular THC norm kernel led to optimal initial conditions that correspond to a vanishing THC index, as is the case here to a good approximation. In those simpler, zonally averaged models (Tziperman and Ioannou 2002; Zanna and Tziperman 2005), the THC was related to the meridional density gradients in some very simple way, following the usual Stommel (1961) formulation. This means that the density gradients also vanish or satisfy some simple condition that leads to a vanishing initial THC. However, in the coupled GCM analyzed here, the THC is clearly not that simply related to the density field, as

the optimal initial state for the density, shown in Figs. 11 and 12, is characterized by nontrivial density gradients. That the initial THC anomaly corresponding to the optimal initial conditions for the temperature and salinity is still very small, as implied from the large amplification factors, is seen in Fig. 9.

As in simpler models, our analysis is able to find initial conditions for the temperature and salinity anomaly for which the THC nearly vanishes, but which lead to a growth at later times due to the different evolutions of the temperature and salinity anomalies. It is encouraging to see the nonnormal amplification operating both in the GCM and in the simpler models. However, the above description of the optimal initial conditions and maximum amplified state is unsatisfactory, as we cannot explain which physical process act to create the amplification as was possible using the simpler models. One may hope to be able to repeat the analysis here using the actual linearized GCM equations, which will enable a more explicit and satisfactory physical interpretation.

One thing that we can conclude from our analysis does shed some new light on the dynamics of THC variability. The small amplitude of present-day THC variability in GCMs seems to suggest that the dynamics of this variability are mostly linear. Griffies and Tziperman (1995) and others suggested that the THC variability may be described as the excitation of a damped oscillatory mode of the THC by stochastic atmospheric forcing. Our reduced space model does, indeed, have several damped and several damped oscillatory modes. The typical decay and oscillation time scales vary from decades to hundreds of years. However, we find here that the dynamics of transient amplification of the THC involve the interaction of more than a single mode (such interaction is necessary for transient amplification to occur). Thus the picture of how THC variability is excited by the atmosphere may be more complex than previously thought.

4. Conclusions

We investigated the amplification of temperature, salinity, and THC anomalies, in the GFDL coupled atmosphere–ocean general circulation model CM2.1, due to nonnormal dynamics. The analysis was done by first constructing a reduced space based on EOF decomposition of temperature and salinity anomalies in the North Atlantic. Instead of describing the evolution of the three-dimensional temperature and salinity fields, the problem was reduced to the evolution of the amplitudes of some 25 principal components (EOF amplitudes) and then to the fitting of a linear model (propa-

gator matrix) to the evolution of these principal components.

By analyzing the resulting linear model, we found that the nonnormal dynamics can lead to a significant growth of temperature, salinity, and THC anomalies with a typical growth time of 5–10 yr. This result leads to a possible order of magnitude estimate of predictability time for the THC in the model. Our analysis was limited by the length of the control integration used to calculate the fitted linearized model (2000 yr). It seems that a longer integration should allow the use of more EOF amplitudes in the reduced state reconstruction, and our results indicate that this will likely result in even stronger transient amplification. We therefore view our results as a lower bound on the nonnormality of THC dynamics in this GCM. This seems to suggest that such nonnormal effects should be seriously considered if the predictability of the THC is to be quantitatively evaluated from models or observations. We should mention that the results found in this work might be different if we were to consider the entire Atlantic basin instead of only the North Atlantic, as suggested by some numerical GCM experiments.

The suggested importance of nonnormal THC dynamics to THC predictability means that one may need to use ensemble runs for THC prediction, following the usual practice in numerical weather prediction. These predictions would need to be initialized by the singular vectors of the linearized dynamics in order to obtain a reliable estimate of the prediction uncertainty. Owing to the nonnormality of the linearized dynamics, the structure of the initial conditions will influence the amplification of the THC after a given amount of time and one cannot arbitrarily start from some state of the ocean–atmosphere and expect to estimate reliable predictability limits.

Griffies and Tziperman (1995) and others suggested that the present-day variability of the THC may be viewed as the excitation of a damped oscillatory mode of the ocean by atmospheric stochastic forcing. The results of this work seem to put this proposal in a different perspective. While stochastic atmospheric excitation is still likely to be important, it seems that, due to the nonnormality of the linearized dynamics, the interaction of several nonorthogonal ocean modes—rather than a single damped oscillatory mode—may be playing a very important role in amplifying the response to the atmospheric stochastic forcing. While it is unfortunate that this picture is not as simple as that of a single mode excited by the atmosphere, this does open up the possibility to explore a richer dynamics using both models and observations.

The need for a long control integration for evaluating

the linear fit and nonnormal effects makes our results tentative. It will be useful to repeat this analysis by using a direct representation of the linearized dynamics of the GCM, which implies using both the model and its adjoint. This will also enable the elucidation of the precise physical mechanisms of the transient growth, something that was not possible in the approach taken here. Such a fuller GCM-based analysis is a technically complex task, but one that is not completely out of reach in the next few years. It is also important to keep in mind that nonlinear processes that we have parameterized as a stochastic process may be playing an important role in the THC oscillations in the GCM, and further analysis is needed to quantify these effects.

We showed that the optimal initial conditions made physical sense (e.g., in terms of their dependence on the weighting by grid element volume and standard deviation of the variability used for nondimensionalizing the temperature and salinity). This suggests that, even before using a full GCM and its adjoint to repeat the analysis done here, an intermediate future step could perhaps be to use the results of the present analysis to initialize a GCM. The GCM could then be used to investigate the physical growth mechanisms and the predictability implications of the nonnormal dynamics found here. It is common practice to evaluate THC predictability by running ensemble prediction experiments started by perturbing only the atmospheric state, using identical ocean states (e.g., Griffies and Bryan 1997; Collins et al. 2006). This has the advantage of evaluating, in a sense, the “best case scenario” predictability. Instead, perturbing the initial ocean state using the methodology presented here may provide a stricter, hence perhaps more realistic, estimate of the predictability.

Acknowledgments. It is a pleasure to thank Brian Farrell, Lenny Smith, Peter Huybers, and Andrew Moore for very useful discussions and advice, and to Steve Griffies and Ed Hawkins for their helpful reviews. We are grateful to Keith Dixon and Tom Delworth for providing the GFDL model output used for this work. This work was supported by the U.S. National Science Foundation Climate Dynamics Program Grant ATM-0351123 and by the McDonnell Foundation. ET thanks the Weizmann Institute, where some of this work was completed during a sabbatical, for its hospitality.

REFERENCES

- Bjerknes, J., 1964: Atlantic air–sea interaction. *Advances in Geophysics*, Vol. 10, Academic Press, 1–82.

- Blumenthal, M. B., 1991: Predictability of a coupled ocean-atmosphere model. *J. Climate*, **4**, 766–784.
- Collins, M., and Coauthors, 2006: Interannual to decadal climate predictability in the North Atlantic: A multimodel-ensemble study. *J. Climate*, **19**, 1195–1203.
- Delworth, T., 2006: Preface. *J. Climate*, **19**, 641.
- , and R. Zhang, 2007: Simulated interdecadal variability of the Atlantic THC in the GFDL CM2.1 climate model. *Geophysical Research Abstracts*, Vol. 9, Abstract 10075. [Available online at <http://www.cosis.net/abstracts/EGU2007/10075/EGU2007-A-10075-2.pdf>.]
- , S. Manabe, and R. J. Stouffer, 1993: Interdecadal variations of the thermohaline circulation in a coupled ocean-atmosphere model. *J. Climate*, **6**, 1993–2011.
- , and Coauthors, 2006: GFDL's CM2 global coupled climate models. Part I: Formulation and simulation characteristics. *J. Climate*, **19**, 643–674.
- Farrell, B., 1988: Optimal excitation of neutral Rossby waves. *J. Atmos. Sci.*, **45**, 163–172.
- , 1989: Optimal excitation of baroclinic waves. *J. Atmos. Sci.*, **46**, 1193–1206.
- , and P. J. Ioannou, 1996: Generalized stability theory. Part I: Autonomous operators. *J. Atmos. Sci.*, **53**, 2025–2040.
- , and —, 2001: Accurate low-dimensional approximation of the linear dynamics of fluid flow. *J. Atmos. Sci.*, **58**, 2771–2789.
- Gnanadesikan, A., and Coauthors, 2006: GFDL's CM2 global coupled climate models. Part II: The baseline ocean simulation. *J. Climate*, **19**, 675–697.
- Griffies, S. M., and E. Tziperman, 1995: A linear thermohaline oscillator driven by stochastic atmospheric forcing. *J. Climate*, **8**, 2440–2453.
- , and K. Bryan, 1997: Predictability of North Atlantic multidecadal climate variability. *Science*, **275**, 181–184.
- , and Coauthors, 2005: Formulation of an ocean model for global climate simulations. *Ocean Sci.*, **1**, 45–79.
- Hasselmann, K., 1988: PIPs and POPs: The reduction of complex dynamical systems using principal interaction and oscillation patterns. *J. Geophys. Res.*, **93**, 11 015–11 021.
- Hawkins, E., and R. Sutton, 2007: Variability of the Atlantic thermohaline circulation described by three-dimensional empirical orthogonal functions. *Climate Dyn.*, **29**, 745–762.
- Jungclauss, J. H., H. Haak, M. Latif, and U. Mikolajewicz, 2005: Arctic-North Atlantic interactions and multidecadal variability of the meridional overturning circulation. *J. Climate*, **18**, 4013–4031.
- Kleeman, R., and A. M. Moore, 1997: A theory for the limitation of ENSO predictability due to stochastic atmospheric transients. *J. Atmos. Sci.*, **54**, 753–767.
- Kushnir, Y., 1994: Interdecadal variations in North Atlantic sea surface temperature and associated atmospheric conditions. *J. Climate*, **7**, 141–157.
- 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*, **51A**, 326–336.
- Moore, A. M., and R. Kleeman, 2001: The differences between the optimal perturbations of coupled models of ENSO. *J. Climate*, **14**, 138–163.
- Penland, C., 1989: Random forcing and forecasting using principal oscillation pattern analysis. *Mon. Wea. Rev.*, **117**, 2165–2185.
- , 1996: A stochastic model of IndoPacific sea surface temperature anomalies. *Physica D*, **98**, 534–558.
- , and P. D. Sardeshmukh, 1995: The optimal growth of tropical sea surface temperature anomalies. *J. Climate*, **8**, 1999–2024.
- , and L. Matrosova, 1998: Prediction of tropical Atlantic sea surface temperatures using linear inverse modeling. *J. Climate*, **11**, 483–496.
- , and —, 2001: Expected and actual errors of linear inverse model forecasts. *Mon. Wea. Rev.*, **129**, 1740–1745.
- , M. Flügel, and P. Chang, 2000: Identification of dynamical regimes in an intermediate coupled ocean-atmosphere model. *J. Climate*, **13**, 2105–2115.
- Stommel, H., 1961: Thermohaline convection with two stable regimes of flow. *Tellus*, **13**, 224–230.
- Stouffer, R. J., and Coauthors, 2006: GFDL's CM2 global coupled climate models. Part IV: Idealized climate response. *J. Climate*, **19**, 723–740.
- Tziperman, E., and P. J. Ioannou, 2002: Transient growth and optimal excitation of thermohaline variability. *J. Phys. Oceanogr.*, **32**, 3427–3435.
- Von Storch, H., T. Bruns, I. Fischer-Bruns, and K. Hasselmann, 1988: Principal oscillation pattern analysis of the 30–60 day oscillation in a GCM equatorial troposphere. *J. Geophys. Res.*, **93**, 11 021–11 036.
- Wittenberg, A. T., A. Rosati, N.-C. Lau, and J. J. Ploshay, 2006: GFDL's CM2 global coupled climate models. Part III: Tropical Pacific climate and ENSO. *J. Climate*, **19**, 698–722.
- Zanna, L., and E. Tziperman, 2005: Nonnormal amplification of the thermohaline circulation. *J. Phys. Oceanogr.*, **35**, 1593–1605.
- , and —, 2008: Optimal surface excitation of the thermohaline circulation. *J. Phys. Oceanogr.*, in press.