The Atmospheric Response over the North Atlantic to Decadal Changes in Sea Surface Temperature

S. VENZKE

Max-Planck-Institut für Meteorologie, Hamburg, Germany

M. R. ALLEN

Atmospheric, Oceanic and Planetary Physics, University of Oxford, Oxford, United Kingdom, and Space Science Department, Rutherford Appleton Laboratory, Chilton, United Kingdom

R. T. SUTTON

Atmospheric, Oceanic and Planetary Physics, University of Oxford, Oxford, United Kingdom, and Department of Meteorology, University of Reading, Reading, United Kingdom

D. P. ROWELL

Hadley Centre for Climate Prediction and Research, Meteorological Office, Bracknell, United Kingdom

(Manuscript received 12 May 1998, in final form 5 November 1998)

ABSTRACT

Decadal fluctuations in the climate of the North Atlantic–European region may be influenced by interactions between the atmosphere and the Atlantic Ocean, possibly as part of a coupled ocean–atmosphere mode of variability. For such a mode to exist, a consistent atmospheric response to fluctuations in North Atlantic sea surface temperatures (SST) is required. Furthermore, this response must provide feedbacks to the ocean. Whether a consistent response exists, and whether it yields the required feedbacks, are issues that remain controversial. Here, these issues are addressed using a novel approach to analyze an ensemble of six integrations of the Hadley Centre atmospheric general circulation model HadAM1, all forced with observed global SSTs and sea-ice extents for the period 1949–93.

Characterizing the forced atmospheric response is complicated by the presence of internal variability. A generalization of principal component analysis is used to estimate the common forced response given the knowledge of internal variability provided by the ensemble. In the North Atlantic region a remote atmospheric response to El Niño–Southern Oscillation and a further response related to a tripole pattern in North Atlantic SST are identified. The latter, which is most consistent in spring, involves atmospheric circulation changes over the entire region, including a dipole pattern in sea level pressure often associated with the North Atlantic oscillation. Only over the tropical/subtropical Atlantic, however, does it account for a substantial fraction of the total variance. How the atmospheric response could feed back to affect the ocean, and in particular the SST tripole, is investigated. Several potential feedbacks are identified but it has to be concluded that, because of their marginal consistency between ensemble members, a coupled mode that relied on these feedbacks would be susceptible to disruption by internal atmospheric variability. Notwithstanding this conclusion, the authors' results suggest that predictions of SST evolution could be exploited to predict some aspects of atmospheric variability over the North Atlantic, including fluctuations in spring of the subtropical trade winds and the higher latitude westerlies.

1. Introduction

There has been much recent interest in the nature and mechanisms of long-term climate variability in the North Atlantic region. This interest has largely been stimulated by the belief that fluctuations on seasonal to decadal timescales might be predictable. Such predictability could arise from the influence of ocean dynamics on sea surface temperatures (SST) and the subsequent influence of SST changes on the atmosphere. However, the mechanisms giving rise to decadal climate variations are poorly understood.

Studies of the observed characteristics of decadal climate variability in the North Atlantic region have identified a weak fluctuation with a period of 12–14 yr iden-

Corresponding author address: Dr. Stephan Venzke, Max-Planck-Institut für Meteorologie, Bundesstraβe 55, D-20146 Hamburg, Germany.
E-mail: venzke@dkrz.de

tifiable in both oceanic and atmospheric variables (Deser and Blackmon 1993; Hansen and Bezdek 1996; Sutton and Allen 1997). The existence of slowly propagating SST anomalies suggests a significant role for ocean dynamics, while local correlations between surface atmospheric variables and SST suggest that ocean mixed-layer processes are also involved. The key issues that remain unclear are whether the changes in SST are having any significant effect on the atmosphere, and what is determining the timescale of the observed fluctuations.

The impact of external forcing due to solar variability, volcanic activity, and anthropogenic emissions has been discussed in several publications (e.g., White et al. 1997; Lean et al. 1995; Robock and Mao 1995) as a possible source of long-term climate fluctuations. Interannual variations of tropical Pacific SSTs associated with the El Niño–Southern Oscillation (ENSO) phenomenon were also found to have a remote effect on the atmospheric circulation over the Atlantic region (e.g., Davies et al. 1997). While these external influences no doubt have some role, much of the character of decadal climate fluctuations is likely to be determined by the Atlantic ocean–atmosphere system itself. There are two general scenarios for this internal variability.

In the first, which may loosely be termed ''uncoupled,'' the atmospheric variability is viewed as essentially stochastic (possessing a ''white'' spectrum) and is insensitive to changes in SST. Frankingoul et al. (1997) have recently generalized Hasselmann's stochastic climate model (Hasselmann 1976) by incorporating ocean dynamics as well as mixed-layer processes. Saravanan and McWilliams (1998) have furthermore described how the ocean dynamics can provide a velocity scale, which, in concert with a spatial scale set by the atmospheric forcing, could lead to a prefered timescale for decadal fluctuations in the ocean. There could be a weak feedback onto the atmosphere but this is not necessary to sustain an oscillation in the ocean. One should also note in this context that the atmosphere might also be able to create some low-frequency variability on its own by internal nonlinear processes (James and James 1989) or through interactions with the land cryosphere. Whether such variability can survive external forcing by the seasonal cycle remains open to question (Lorenz 1990).

In the second, coupled, scenario a consistent atmospheric response to Atlantic SST anomalies is a requirement. Latif and Barnett (1994) proposed a mechanism for decadal variability in the North Pacific in which the atmospheric response involved: first, surface heat flux anomalies producing a positive feedback on SST anomalies, and second, anomalies in wind stress curl producing a delayed negative feedback via adjustment of the subtropical gyre circulation, with the delay time set by the propagation of baroclinic Rossby waves across the ocean basin. Grötzner et al. (1998), studying the Atlantic with the same coupled model integration as Latif and Barnett (1994), interpreted their results in terms of this mechanism. In so doing, they noted that advection in the ocean could be as important as Rossby wave propagation in providing the delayed negative feedback.

The decisive factor that can discriminate between the uncoupled and coupled scenarios is the nature of the atmospheric response to Atlantic SST anomalies. The coupled scenario places stringent requirements on the characteristics of the response. First, it must be sufficiently strong that its influence on the ocean is not dominated by the influence of internal atmospheric variability. Second, it must involve changes in surface wind stress and/or heat flux that have an appropriate form to provide a feedback that will modulate, with a lag, the SST features to which the atmosphere is responding.

The nature of the atmospheric response to various idealized extratropical SST anomalies has already been investigated in several model studies (Palmer and Sun 1985; Pitcher et al. 1988; Kushnir and Held 1996; Peng et al. 1997). The results from these studies are not consistent in all respects, but tend to suggest that the response is rather weak relative to internal variability. This fact immediately presents a signal-to-noise problem: to identify the atmospheric response reliably we need information about the internal variability. Such information is available from an ensemble of model integrations.

An analysis of signal-to-noise in the ensemble of model integrations that we analyze was carried out by Davies et al. (1997). Their approach allows regions where a forced response may exist to be identified, but does not directly allow the delineation of the spatio-temporal characteristics of the forced response. Allen and Smith (1997) recently presented a generalization of principal component (or singular spectrum) analysis (PCA or SSA), which yields an optimal multivariate filter to discriminate between a signal and some arbitrary colored noise. Their ideas are related to those of ''optimal fingerprinting'' (Hasselmann 1979, 1997) and a similar method was presented by Thacker and Lewandowicz (1996) for pattern detection in the presence of correlated analysis errors.

In this paper we use an ensemble of multidecadal atmospheric GCM runs, forced with observed SSTs, in conjunction with this optimal detection algorithm, to study the atmospheric response to observed decadal SST changes. In section 2 we document the model, the experimental design, and the data treatment. Section 3 contains the derivation of the statistical method used. In section 4 we isolate the effect of North Atlantic SST anomalies, discuss the characteristics of the obtained reponse, compare it with observations, and report the sensitivity of the statistical method used. We conclude the paper with a summary of our major findings and their implications in section 5.

2. Experimental design and data treatment

The GCM used in this paper is known as HadAM1, the version of the Hadley Centre Atmosphere Model

that was submitted to the Atmospheric Model Intercomparison Project (Gates 1992). It is a gridpoint model with a horizontal resolution of 2.5° lat × 3.75° long and 19 hybrid levels in the vertical. Physical parameterizations include: a gravity wave drag scheme; a radiation scheme that computes fluxes in four longwave bands and six shortwave bands and responds to prognostic cloud variables; a penetrative convection scheme with stability-dependent closure; and a land surface scheme. Further details about the model development and the physics and dynamics of the precise version used here are given by Cullen (1993) and Phillips (1994), respectively.

The simulations used here are the same six-member ensemble employed by Rowell (1998) and Davies et al. (1997). Each member was integrated from 1 October 1948 to 1 December 1993 with lower boundary forcing supplied by version 1.1 of the specifically designed Global Sea-Ice and Sea Surface Temperature (GISST) dataset, described by Parker et al. (1995). The initial atmospheric state and surface temperatures were taken from U.K. Meteorological Office analyses for six different dates, arbitarily chosen, but all close to 1 October. Soil moisture and snow depth were initialized from an adaption of the Willmott et al. (1985) climatology. Due to likely spinup effects, the first two months of each integration were discarded, leaving a set of six integrations, each 45-yr long (1 December 1948–30 November 1993), giving a total of 270 yr of model data.

Whether the model's response to SST is realistic or not depends in part on the accuracy of various aspects of its climatology. The climatologies of model variables of interest in this paper were computed for each month by averaging that month over all 270 model years. As shown by Davies et al. (1997), the seasonal ensemble mean climatology of mean sea level pressure (MSLP) over the North Atlantic–European area compared very well with that inferred from the observational dataset of Basnett and Parker (1997).

Monthly anomalies were obtained by subtracting the climatology from the monthly data. Standard seasonal means [December–February (DJF), March–May (MAM), June–August (JJA), September–November (SON)] were constructed by averaging over 3-month periods, and each season was analyzed separately. Since we focus in this study on low-frequency atmospheric changes, we emphasized this timescale by removing any linear trend and smoothing each set of seasonal means with a three-point (i.e., 3-yr) running mean filter.

In this paper we will focus on results for winter (DJF) and spring (MAM) data over the North Atlantic. Most studies of coupled air–sea modes in midlatitudes restrict their analyses to wintertime when the oceanic mixed layer is deepest and SST anomalies are most likely to express deeper subsurface anomalies (e.g., Deser and Blackmon 1993; Grötzner et al. 1998; Sutton and Allen 1997; Latif and Barnett 1994). Studies of seasonal predictability over the North Atlantic and European region by Davies et al. (1997) as well as Palmer and Anderson (1994), however, have suggested that the MSLP predictability over Europe may be higher in spring (MAM) than in winter (DJF).

3. Detection of SST-forced atmospheric variability

In this section we present a technique for characterizing the time-varying response to a time-varying external forcing in a system that is subject to internal chaotic variability, making use of the fact that we have an ensemble of realizations with identical external forcing (SST and sea ice) at our disposal. The definition of the time-mean response to a constant external forcing is relatively straightforward, but the definition of the time-varying response is somewhat more problematic. For the sake of simplicity, we assume an estimate of the 45-yr-mean has been subtracted from all the quantities considered, and confine our attention to the ''first-order'' time-varying forced response, referred to hereafter as simply ''the forced response.'' This we define as the component of the evolution of the system that is determined by the external forcing, independent of the initial conditions, and would therefore be common to all members of a hypothetical infinite ensemble. If we had such an infinite ensemble available, the forced response would simply be the time-varying component of the ensemble mean: the objective of this section is to estimate this response making the most efficient use possible of the information provided by a relatively small (six member) ensemble that displays substantial intermember variability.

Since we are dealing with seasonal means, sampled annually, the ''atmospheric noise'' (internal variability that is independent of SST and sea ice forcing) is uncorrelated in time: each season is effectively independent of the same season the following year. This noise is, however, highly correlated in space—that is, certain spatial patterns [such as that associated with the North Atlantic Oscillation (NAO; Hurrell 1995)] contain much more internal variability than others. Assuming that some forced response exists [already established by Davies et al. (1997)], we require three diagnostics to describe its basic spatio-temporal characteristics: (i) the spatial pattern, which provides an optimal filter to discriminate between the forced response and internal variability (required to compute the other diagnostics); (ii) the time series characterizing the temporal evolution of the dominant mode of the forced response (obtained by projecting the data onto this optimal spatial filter); and (iii) the pattern characterizing the spatial characteristics of the dominant mode of the forced response (obtained by projecting the data onto this time series).

We stress that not all possible atmospheric responses to SST and sea-ice forcing would be captured by our definition of the forced response. We assume that, on the timescales of interest, the data may be treated as a linear superposition of a deterministic and predictable

forced response with additive internal variability. This picture is clearly incomplete, since in a nonlinear system the internal variability and forced response will interact. For example, the response to a particular pattern of SST anomalies could be enhanced variability in the storm track accompanied by little change in the mean flow. This kind of higher-order response would not systematically influence a diagnostic based on ensemble means. Focusing on the first-order response seems a logical starting point, and because we are working with averages over timescales, which are long relative to those of the components of the system which are known to be highly chaotic, the linearity approximation may be acceptable (Allen et al. 1997).

In the language of weather regimes (e.g., Lorenz 1963; Robertson et al. 1999, manuscript submitted to *J. Atmos. Sci.*; Renshaw et al. 1998), the first-order forced response corresponds to either (a) changes in the characteristics of the regimes themselves, or perhaps more likely (b) changes in the fraction of time the system spends in particular regimes, in response to changes in the state (or ''phase'') of the external forcing. A higher-order response would correspond to changes in the forcing introducing systematic changes in the frequency of regime transitions. Note that our definition of the forced response is applicable regardless of whether the response consists in the excitation of a preexisting mode of internal variability or the generation of a new mode.

An important assumption in our approach is that all distributions are Gaussian. Other approaches (e.g., cluster analysis) might be more effective for highly non-Gaussian systems, but since none of the quantities considered here show significant departures from Gaussian behavior on the timescales of interest (Rowell 1998), we believe our procedure is appropriate.

Readers who feel they now have a sufficient grasp of the principle and are uninterested in the estimation procedure itself may wish to proceed directly to our results in section 4.

a. Summary of the analysis procedure

The analysis procedure we use is detailed in Allen and Smith (1997) and is based on the optimal fingerprinting algorithm of Hasselmann (1979), generalized to the situation in which we do not know the spatio-temporal characteristics of the response pattern we are looking for. Similar ideas for different applications may be found in Zhang et al. (1994) and Thacker and Lewandowicz (1996).

We represent the data in the kth ensemble member as the matrix $\mathbf{X}_k$, where X_{ijk} is the DJF or MAM seasonal mean at spatial location i ($i = 1, l$) in year j ($j = 1, m$) from ensemble member k ($k = 1, n$). The seasons are analyzed independently, so m is the number of years in the series. The best estimate of the 45-yr-mean component of the response, obtained by averaging over time and over the ensemble, is subtracted from the data at each spatial location prior to the analysis. This needs to be taken into account in the formulation of statistical tests, but for brevity we shall not mention it again in the discussion.

To estimate the true forced response, $\mathbf{X}_F$, from the finite ensemble, we decompose $\mathbf{X}_k$ into an ensemble mean component $\mathbf{X}_M$ and the departures from that mean, $\mathbf{X}_{Nk}$:

$$\mathbf{X}_k = \mathbf{X}_M + \mathbf{X}_{Nk}. \tag{1}$$

By definition, the forced response discussed above, $\mathbf{X}_F = \mathcal{E}(\mathbf{X}_M)$, where the expectation operator refers to an average taken over an hypothetical infinite ensemble, not over time. Here $\mathbf{X}_F$ is therefore time dependent, having the same dimensions as the $\mathbf{X}_k$. The ensemble mean, $\mathbf{X}_M$, is the best estimate we have of the full spatio-temporal evolution of the forced response, but it will be heavily contaminated by noise due to the small size of the ensemble. The aim is to identify the dominant characteristics of $\mathbf{X}_F$ despite this contamination. Our procedure may be summarized as follows.

First, we append the six fields $\mathbf{X}_{Nk}$ to a set of $6 \times 45 = 270$ fields ($\mathbf{X}_N$) and perform a singular value decomposition (SVD) on this set, that is:

$$\mathbf{X}_N = \mathbf{E}_N \mathbf{\Lambda}_N \mathbf{P}_N^{\mathrm{T}}, \tag{2}$$

ranking the elements of $\mathbf{\Lambda}_N$, as conventional, in decreasing order.

Second, we project the ensemble mean $\mathbf{X}_M$ onto the κ highest-ranked ''EOFs'' of the noise (i.e., the first κ columns of the matrix $\mathbf{E}_N$, being the left singular vectors of $\mathbf{X}_{Nk}$), weighting by the inverse of the corresponding noise singular values:

$$\mathbf{X}'_M = \mathbf{F}^{\mathrm{T}} \mathbf{X}_M = n^{1/2} (\mathbf{\Lambda}_N^{(\kappa)})^{-1} \mathbf{E}_N^{(\kappa)\,\mathrm{T}} \mathbf{X}_M. \tag{3}$$

The effect of the ''prewhitening operator'' $\mathbf{F}$ is to give all patterns of internal variability contaminating $\mathbf{X}_M$ the same expected variance, or weight, in $\mathbf{X}'_M$. The factor of $n^{1/2}$ ensures that, in the absence of any true forced response, $\mathbf{X}'_M$ will appear as unit-variance white noise.

Third, we perform an SVD on the prewhitened ensemble mean:

$$\mathbf{X}'_M = \mathbf{E}' \mathbf{\Lambda}' \mathbf{P}^{\mathrm{T}}. \tag{4}$$

The convolution of $\mathbf{E}'$ with the prewhitening operator, $\tilde{\mathbf{E}} = \mathbf{F}\mathbf{E}'$, gives the set of optimal spatial filters we are looking for. The highest ranked filter, $\tilde{\mathbf{e}}_1$, is the spatial pattern that maximizes the ratio of ensemble-mean to within-ensemble variance, that is, $\tilde{\mathbf{e}}_1^{\mathrm{T}} \mathbf{X}_M \mathbf{X}_M^{\mathrm{T}} \tilde{\mathbf{e}}_1 / \sum_k \tilde{\mathbf{e}}_1^{\mathrm{T}} \mathbf{X}_{Nk} \mathbf{X}_{Nk}^{\mathrm{T}} \tilde{\mathbf{e}}_1$: see Eqs. (15) and (16) below. The time series obtained by projecting $\mathbf{X}_M$ onto $\tilde{\mathbf{e}}_1$ and normalizing,

$$\mathbf{p}_1 = \frac{\mathbf{X}_M^{\mathrm{T}} \tilde{\mathbf{e}}_1}{\lambda'_1}, \tag{5}$$

which is also highest-ranked right singular vector of $\mathbf{X}'_M$, characterizes the temporal evolution of the dominant mode of the forced response.

Fourth, we characterize the spatial pattern associated

with the dominant forced response by projecting $\mathbf{X}_M$ onto this ''optimized: PC''

$$\hat{\mathbf{e}}_1 \simeq \frac{\mathbf{X}_M \mathbf{p}_1}{\lambda_1'}. \quad (6)$$

This last equality is only approximate, since the exact definition of $\hat{\mathbf{e}}_1$ depends on κ [see Eqs. (27)–(30) below]. A detailed explanation and application to the HadAM1 ensemble is given in the following subsections.

b. Limitations of principal component analysis and analysis of variance

The simplest way of understanding how this optimal filtering procedure works, and why it is necessary, is in the context of the conventional method of analyzing spatio-temporal datasets of this nature, principal component analysis (PCA), which we discuss first. The basic assumption of PCA is that $\mathbf{X}_M$ may be represented by a small number of spatial patterns, or EOFs (empirical orthogonal functions; Lorenz 1956), modulated by corresponding time series of pattern amplitudes, or PCs. We focus on mean sea level pressure (MSLP) in this section (see section 4 for the application to other atmospheric variables).

Figures 1a,b show the leading EOFs and associated normalized PCs of the ensemble mean MSLP over the North Atlantic for winter and spring. They account for 40% and 50% of the total variance in each season, respectively. The PCs suggest variations on a decadal timescale (but recall we are dealing with temporally smoothed data), while the EOF patterns display the dipole structure of the atmospheric part of the decadal mode described by Deser and Blackmon (1993) and Grötzner et al. (1998), which is often associated with the NAO. The seasonal difference between the patterns is similar to that observed by Rogers (1990) and Davies et al. (1997). Differences in detail between the patterns shown here and those of Davies et al. (1997) are due to our use of ensemble means, temporal smoothing to focus on low-frequency variability and a different analysis domain.

These EOF–PC pairs, by construction, account for the maximum possible variance in the ensemble-mean (MSLP), $\mathbf{X}_M$. However, they also account for a substantial fraction of the internal atmospheric variability, whose characteristics may be estimated from the deviations from the ensemble mean, $\mathbf{X}_{Nk}$. As a result, projections of the individual ensemble members onto these EOFs, shown as thin lines in Fig. 1b, display a large spread, raising the question of whether the true forced response in this pattern might be zero, and the ensemble-mean variability, which we observe (thick line in Fig. 1b), is simply due to sampling noise in a finite ensemble.

This question can be addressed straightforwardly, using the analysis of variance (ANOVA) approach of Davies et al. (1997) and Rowell (1998). If the true forced response in this pattern is zero and the ensemble members are independent and normally distributed,[1] then the expected variance of the ensemble mean time series σ_M^2 would be $1/n$ times the expected variance of the internal variability σ_N^2. This can be couched in terms of a statistical test, as follows: if $\mathbf{x}_k$ is the time series of amplitude coefficients of a normalized pattern $\mathbf{e}$ (the first EOF of the ensemble mean in this case) in the kth ensemble member, so $\mathbf{x}_k = \mathbf{X}_k^{\mathrm{T}}\mathbf{e}$, then

$$n\frac{\hat{\sigma}_M^2}{\hat{\sigma}_N^2} \equiv n\frac{\dfrac{1}{m-1}\mathbf{X}_M^{\mathrm{T}}\mathbf{X}_M}{\dfrac{1}{m(n-1)}\displaystyle\sum_k \mathbf{X}_{Nk}^{\mathrm{T}}\mathbf{X}_{Nk}} \sim F_{m'-1,m'(n-1)}, \quad (7)$$

where $F_{m'-1,m'(n-1)}$ is the standard F-distribution with $m' = m/3$ to take into account the 3-yr running mean smoothing, and $\hat{\sigma}_M^2$ and $\hat{\sigma}_N^2$ are unbiased estimates of σ_M^2 and σ_N^2, respectively [the factors of $m - 1$ and $n - 1$ in their definitions take into account the fact that $\Sigma_j (X_M)_{ij} \equiv 0$ and $\Sigma_k (X_N)_{ijk} \equiv 0$ for all i and ij, respectively (see Rowell 1998)].

There is some room for ambiguity in m' on the rhs of Eq. (7) because of our use of temporally smoothed data. We have used a conservative (small) estimate, but none of our conclusions are sensitive to increasing it by up to a factor of two. The top two lines of results in Table 1 show $\hat{\sigma}_M^2$ and $\hat{\sigma}_N^2$, their ratio, and n times their ratio for the projection of the ensemble members onto the first EOF of the ensemble mean as shown in Fig. 1b. Since $n\hat{\sigma}_M^2/\hat{\sigma}_N^2$ exceeds the 95% critical value of the F-distribution, $F_{0.05}$, we can conclude at this confidence level that there is a nonzero forced response in this pattern (or ''state space direction'').

The spread of the thin lines in Fig. 1b, however, indicates that even though this response is nonzero, it is far from robust. It helps to recall at this point what a significant result from such an ANOVA implies in physical terms: it means that the signal-to-noise ratio is such that the time series of the amplitude of the pattern $\mathbf{e}$ in the ensemble mean is likely (>95% chance) to be positively correlated with the time series of this pattern amplitude in the true (infinite ensemble) forced response (we shall refer to this condition as the ''temporal evolution'' of this pattern being of the correct sign). This only means there is a better than 95% chance that the ensemble-mean response will be of the correct sign in significantly over half the individual years (how much over half depends on the temporal characteristics of the

[1] The assumption that the ensemble members are independent in the absence of any forced response to SST and sea ice is quite secure, since the predictability horizon of the atmosphere model is such that initial conditions are unlikely to have any significant influence on the statistics of a 45-yr integration. Rowell (1998) finds no evidence of departures from normality on the timescales of interest here. Evidence of multimodal, nonnormal distributions of atmospheric variables mainly arises on shorter timescales, when examining daily or 5-day averages.

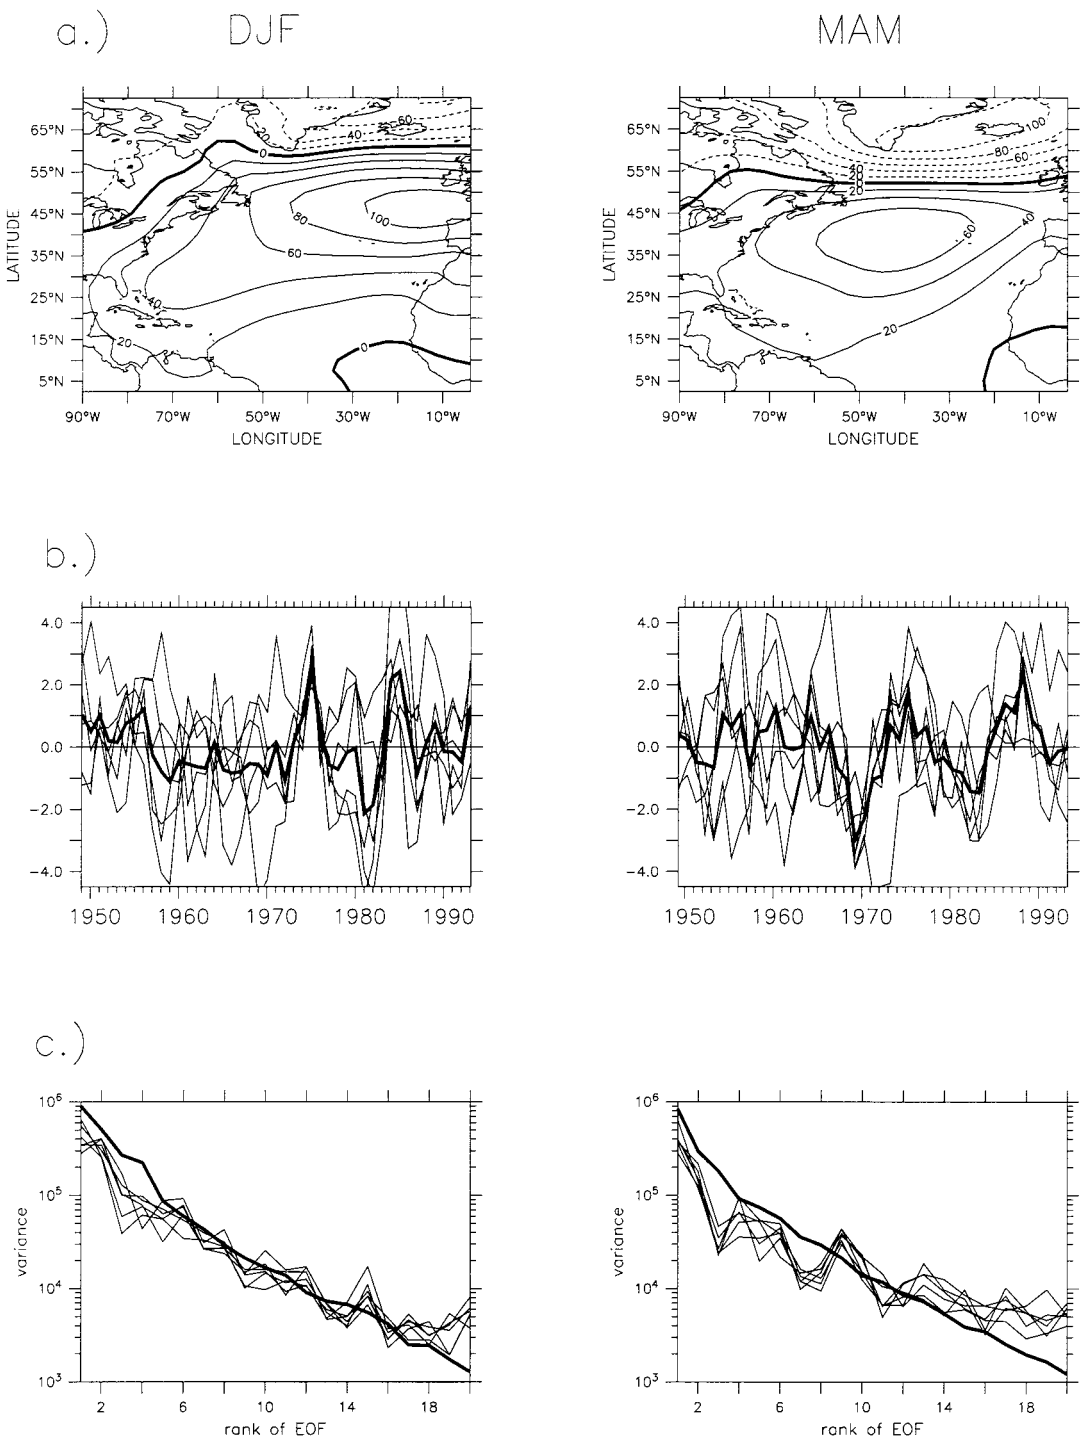

FIG. 1. (a) First EOF of winter and spring MSLP (Pa) for the ensemble mean (explaining 40% and 50% of the total variance, respectively—recall the data has been smoothed with a 3-yr running mean). (b) Projections of MSLP from ensemble members (thin lines) and ensemble mean (thick line) onto patterns shown in (a). (c) Variance of projections of MSLP ensemble mean (heavy line) and noise (thin lines) onto leading 20 EOFs of the MSLP ensemble means.

TABLE 1. Estimated ensemble-mean and within-ensemble variances, $\hat{\sigma}_M^2$ and $\hat{\sigma}_N^2$ in selected diagnostics. The 5% cutoff value of the $F_{m'(n-1)}$ distribution, $F_{0.05} = 1.83$. Ratios, shown bold, in column 4, which exceed this cutoff value indicate a response in the corresponding diagnostic, which is detectable at the 95% level. Bold ratios in column 3 indicate a response that is both detectable and consistent at this level. Proj. ≡ projected. Ens. ≡ ensemble.

Diagnostic	$\hat{\sigma}_M^2$	$\hat{\sigma}_N^2$	Ratio	$n\times$ratio
Proj. on ens. mean EOF-1 (DJF, detrended)	9.15×10^5	2.58×10^6	0.354	**2.128**
Proj. on ens. mean EOF-1 (MAM, detrended)	8.58×10^5	2.45×10^6	0.350	**2.101**
Proj. on S/N max EOF-1 (DJF, detrended)	17.558	6.000	**2.926**	**17.558**
Proj. on S/N max EOF-1 (MAM, detrended)	24.184	6.000	**4.031**	**24.184**
Proj. on ens. mean EOF-1 (DJF, no ENSO)	7.48×10^5	2.80×10^6	0.264	1.586
Proj. on ens. mean EOF-1 (MAM, no ENSO)	8.25×10^5	2.39×10^6	0.345	**2.071**
Proj. on S/N max EOF-1 (DJF, no ENSO)	6.658	6.000	1.109	**6.658**
Proj. on S/N max EOF-1 (MAM, no ENSO)	14.918	6.000	**2.486**	**14.918**

signal and noise). It does not even mean this much for the individual ensemble members.

Results from ANOVA are critically dependent on the size of the ensemble. With these signal-to-noise ratios, had we used an ensemble of size 4 we would not have detected a significant response in this NAO-like dipole in either season. It will always be possible to detect a nonzero but arbitrarily weak forced response given a sufficiently large ensemble. Since the real world represents an ensemble of size one, we need to distinguish between a response that is strong enough to be detectable only in an ensemble mean, and a response that is strong enough that we would at least expect with a certain confidence level that the sign of the individual ensemble members will be the same. We will therefore distinguish between a *detectable* response, in which $n\hat{\sigma}_M^2/\hat{\sigma}_N^2 > F_{0.05}$, and a *consistent* response, in which $\hat{\sigma}_M^2/\hat{\sigma}_N^2 > F_{0.05}$. For simplicity, we use the same (ANOVA-based) cutoff values for the two criteria, although a case might be made for using $F_{m'-1,m'}$ for the test of consistency—this would not affect which numbers are highlighted as passing this test in column three of Table 1.

The requirement of consistency is not satisfied by the NAO-like dipole obtained by standard PCA. Thus we should not expect (at the 95% level) to get even the sign of the temporal evolution of this pattern correct in an individual ensemble member. We would conclude, therefore, that even though there is a detectable forced response in this direction, this response is likely to be too weak to play a significant role in any coupled mode of variability.

c. Derivation of the optimal filter

The fact that the highest-ranked EOF/PCs from standard PCA fail to provide evidence for a consistent forced response does not mean that no such response exists. To understand how we can improve on standard PCA to identify the most consistent aspects of the atmospheric response in an ensemble experiment, it helps to recall that the EOFs from PCA are the eigenvectors of the estimated spatial covariance matrix,

$$\hat{\mathbf{C}}_M \equiv \frac{1}{m-1}\mathbf{X}_M\mathbf{X}_M^{\mathrm{T}}. \tag{8}$$

Assuming that the ensemble members are independent and normally distributed, we have

$$\mathcal{E}(\hat{\mathbf{C}}_M) = \mathbf{C}_F + \frac{1}{n}\mathbf{C}_N, \tag{9}$$

where $\mathbf{C}_F$ and $\mathbf{C}_N$ are the true (unknown) spatial covariances of the forced response and internal variability, respectively. In the limit of an infinite ensemble, $n \to \infty$ and $\hat{\mathbf{C}}_M \to \mathbf{C}_F$, but with the ensemble sizes we are considering here, the variance attributable to the forced response may well be $O(1/n)$ times the variance attributable to internal variability, so both terms on the rhs of Eq. (9) will have comparable magnitude, and both will contribute to the eigenvectors of $\hat{\mathbf{C}}_M$.

The assumption (underlying PCA) that the forced response consists of a small number of spatial patterns implies that $\mathbf{C}_F$ will be rank deficient. The eigenvectors of $\hat{\mathbf{C}}_M$, or EOFs, provide an estimate of the eigenvectors of $\mathbf{C}_F$ if and only if the internal variability is uncorrelated in space, or $\mathbf{C}_N = \sigma^2\mathbf{I}$, since adding $\sigma^2\mathbf{I}$ to any matrix simply raises its eigenvalues by σ^2 and does not change its eigenvectors. But the internal variability is clearly not uncorrelated in space: worse still, it contains high variance in precisely the directions that a standard PCA suggests we should be looking for a forced response. The estimated variance of the projections of the $\mathbf{X}_{Nk}$ onto the EOFs of $\mathbf{X}_M$ are shown (multiplied by a factor of $1/n$ to indicate the variance we would expect in the absence of any forced response) as thin lines in Fig. 1c. EOFs that contain high variance in the ensemble mean also contain high variance in the deviations from that mean. Hence the EOFs of the ensemble mean are biased toward the EOFs of the noise and therefore in general will not be aligned, even approximately, with the EOFs of the true forced response (Allen and Smith 1996).

Allen and Smith (1997) resolve this problem by introducing a ''prewhitening'' transformation, $\mathbf{F}$, such that $\mathbf{F}^{\mathrm{T}}\mathbf{C}_N\mathbf{F} \simeq n\mathbf{I}$ (i.e., $\mathbf{F}$ removes spatial correlations, making the noise spatially white—it is already white in time). The true covariance of internal variability, $\mathbf{C}_N$, is un-

known, but an unbiased estimate can be obtained from the deviations from the ensemble mean:

$$\hat{\mathbf{C}}_N \equiv \frac{1}{m(n-1)} \sum_{k=1}^{n} \mathbf{X}_{Nk}\mathbf{X}_{Nk}^{\mathrm{T}}. \tag{10}$$

Note that this is simply the generalization of $\hat{\sigma}_N^2$ to a multivariate process. If we define the columns of $\mathbf{E}_N^{(\kappa)}$ to be κ highest-ranked normalized eigenvectors (EOFs) of $\hat{\mathbf{C}}_N$, and $\mathbf{\Lambda}_N^{(\kappa)}$ the corresponding diagonal matrix of singular values, we define the prewhitening transformation

$$\mathbf{F} \equiv n^{1/2}\mathbf{E}_N^{(\kappa)}(\mathbf{\Lambda}_N^{(\kappa)})^{-1}, \tag{11}$$

and its transposed pseudo-inverse as

$$\mathbf{F}^{(-1)} \equiv n^{-1/2}\mathbf{E}_N^{(\kappa)}\mathbf{\Lambda}_N^{(\kappa)}, \tag{12}$$

so $\mathbf{F}$ and $\mathbf{F}^{(-1)}$ are $l \times \kappa$ matrices. We are obliged to truncate because variance will generally be underestimated in the low-ranked EOFs of internal variability, so while $\mathbf{F}^{\mathrm{T}}\hat{\mathbf{C}}_N\mathbf{F} \equiv n\mathbf{I}$ for all κ by construction, we expect $\mathbf{F}^{\mathrm{T}}\mathbf{C}_N\mathbf{F} \simeq n\mathbf{I}$ only if we confine the analysis to well-sampled directions in state space (meaning, EOFs that contain a realistic level of variance).

To specify the truncation level, κ, at which we avoid unrealistically low-variance EOFs while still retaining a high proportion of the expected signal, we compute the ratio of the variance of $\mathbf{X}_M$ and the $\mathbf{X}_{Nk}$ in the directions defined by the columns of $\mathbf{E}_N$ (Fig. 2a). The cumulative average of this variance ratio, shown in Fig. 2b, stabilizes at about $\kappa = 15$ for winter data, and $\kappa = 19$ for spring. These are the truncations that we will use in the following. The sensitivity of our results to the choice of the truncation is discussed in section 4.

We apply the prewhitening transformation to $\hat{\mathbf{C}}_M$, such that

$$\mathbf{C}'_M \equiv \mathbf{F}^{\mathrm{T}}\hat{\mathbf{C}}_M\mathbf{F} \simeq \mathbf{C}'_F + \mathbf{I} \tag{13}$$

and diagonalize

$$\mathbf{C}'_M \equiv \mathbf{E}'\mathbf{\Lambda}'^2\mathbf{E}^{\mathrm{T}}, \tag{14}$$

arranging the columns of $\mathbf{E}'$, as conventional, in order of decreasing eigenvalue [note that this equivalent to the SVD in Eq. (4)].

Since the transformed internal variability has equal variance in all state-space directions, the highest-ranked eigenvector $\mathbf{e}'_1$ of $\mathbf{C}'_M$ provides an estimate of the highest-ranked eigenvector of $\mathbf{C}'_F$. The elements of $\mathbf{e}'_1$, however, correspond to EOF indices rather than spatial locations, so it aids interpretation to convolve it with the prewhitening transformation, $\mathbf{F}$, giving $\tilde{\mathbf{E}}$, $= \mathbf{F}\mathbf{E}'$. The vector $\tilde{\mathbf{e}}_1$ is the pattern that, for a given truncation, maximizes the ratio of ensemble-mean variance to within-ensemble variance:

$$n\frac{\hat{\sigma}_M^2}{\hat{\sigma}_N^2} = n\frac{\tilde{\mathbf{e}}_1^{\mathrm{T}}\hat{\mathbf{C}}_M\tilde{\mathbf{e}}_1}{\tilde{\mathbf{e}}_1^{\mathrm{T}}\hat{\mathbf{C}}_N\tilde{\mathbf{e}}_1} = \hat{\mathbf{e}}_1^{\mathrm{T}}\hat{\mathbf{C}}_M\tilde{\mathbf{e}}_1 \tag{15}$$

$$= \mathbf{e}_1'^{\mathrm{T}}\mathbf{F}^{\mathrm{T}}\hat{\mathbf{C}}_M\mathbf{F}\mathbf{e}'_1 = \mathbf{e}_1'^{\mathrm{T}}\mathbf{C}'_M\mathbf{e}'_1 = \lambda_1'^2, \tag{16}$$

where $\hat{\sigma}_M^2$ and $\hat{\sigma}_N^2$ are defined as above but with respect to $\tilde{\mathbf{e}}_1$ instead of $\mathbf{e}$. This vector $\tilde{\mathbf{e}}_1$ is the optimal filter for characterizing the forced response (common variability), and is shown (after normalizing for display purposes) for the winter and spring seasons in Fig. 3a.

The optimal filter patterns (Fig. 3a) reveal large-scale coherent regions over the tropical and subtropical Atlantic Ocean whereas in midlatitudes the patterns are broken up into small-scale structures of lower amplitude. Thus the time series of the amplitude of these patterns in any dataset will be dominated by what occurs at low latitudes. This is consistent with the picture provided by a local ANOVA (see Rowell 1998). The advantage of the method presented here over a local ANOVA is that it allows us to begin to delineate the spatio-temporal characteristics of the dominant forced response, in addition to identifying the regions where such a response may exist.

d. Associated time series, or ''optimized PCs''

The normalized pattern-amplitude time series of $\tilde{\mathbf{e}}_1$ in the ensemble mean $\mathbf{X}_M$, that is, $\mathbf{p}_1 = \lambda_1'^{-1}\mathbf{X}_M^{\mathrm{T}}\tilde{\mathbf{e}}_1$, is shown as the thick line in Fig. 3b. The thin lines in Fig. 3b show the projections of the individual ensemble members onto $\tilde{\mathbf{e}}_1$, that is, $\mathbf{p}_k = \lambda_1'^{-1}\mathbf{X}_k^{\mathrm{T}}\tilde{\mathbf{e}}_1$. The algorithm has clearly been extremely successful in reducing the spread between ensemble members: in the direction defined by $\tilde{\mathbf{e}}_1$, we have no problem at all detecting both a detectable and a consistent response, in the sense defined above (Table 1, line 3 and 4). The thick line in Fig. 3c shows the eigenspectrum of $\mathbf{C}'_M$, being n times the estimated ratio of ensemble-mean to within-ensemble variance in the directions defined by the $\tilde{\mathbf{e}}_i$, [see Eq. (15)]. Note that the $\tilde{\mathbf{e}}_i$ are not normalized in the conventional sense, so projecting a dataset onto them provides information about the ''signal-to-noise ratio'' (the variance relative to the variance we would expect if the dataset is pure noise), rather than the absolute variance in the relevant state-space-direction. The thin lines show the corresponding signal-to-noise ratio for the individual $\mathbf{X}_{Nk}$, that is, the diagonal elements of

$$\mathbf{\Lambda}'^2_{Nk} = \frac{1}{mn(n-1)}\tilde{\mathbf{E}}^{\mathrm{T}}\mathbf{X}_{Nk}\mathbf{X}_{Nk}^{\mathrm{T}}\tilde{\mathbf{E}}. \tag{17}$$

The mean of these thin lines (the ''noise-to-noise ratio'') is unity by construction (dotted line).

Note that, if we had reason to believe the noise (internal variability) was self-correlated in time, then we could further improve the estimate of the forced response by prewhitening in time. Since, however, the within-ensemble variability is uncorrelated year-on-year, nothing would be gained from this additional step: hence the asymmetry between space and time in our analysis.

Figure 3d shows the estimated ensemble-mean variance and $1/n$ times the estimated within-ensemble variance in the directions defined by the $\tilde{\mathbf{E}}$. Comparison with Fig. 1c indicates that while the variance of $\mathbf{X}_M$ in

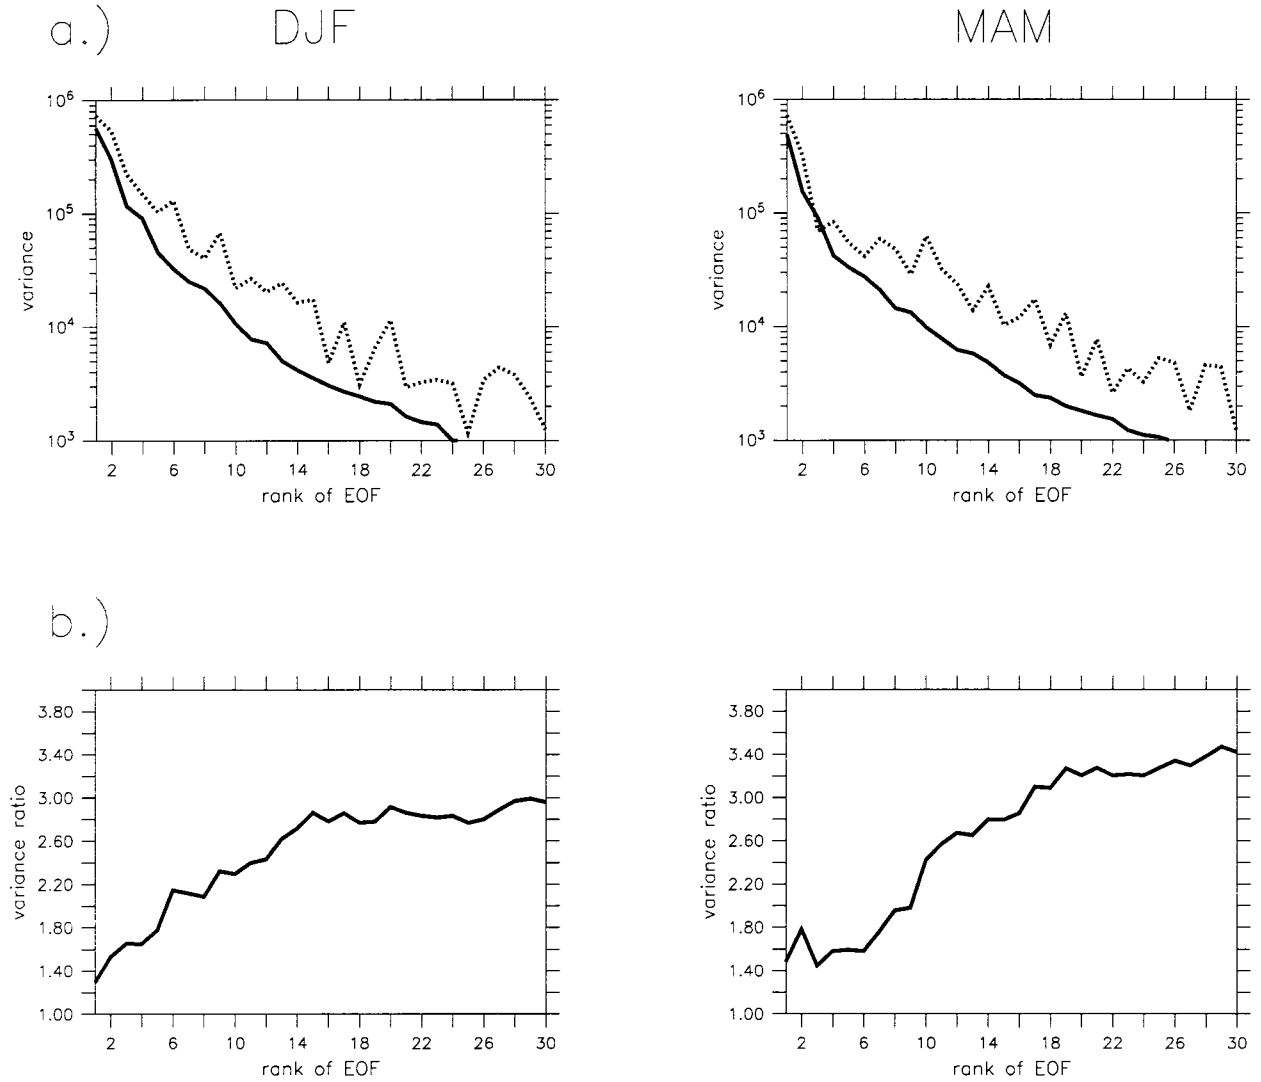

FIG. 2. (a) Variance of projections of ensemble mean $\mathbf{X}_M$ MSLP (dotted line) and deviations from the ensemble mean, $\mathbf{X}_{Nk}$ (solid line) onto leading 30 EOFs of the within-ensemble variability, $\mathbf{E}_N$, in winter and spring. (b) Cumulative averaged ratio $r_\kappa = (1/\kappa)\sum_{i=1}^{\kappa}\sigma_M^2/\sigma_N^2$, with κ = rank of noise EOF, of ensemble-mean variance and within-ensemble variance shown in (a).

the direction of $\tilde{\mathbf{e}}_1$ is a factor of 5.7 (DJF) and 6.8 (MAM) lower than in the direction of the first EOF of $\mathbf{X}_M$, the corresponding variances of the $\mathbf{X}_{Nk}$ have been reduced by about two orders of magnitude, resulting in a significant enhancement of signal-to-noise.

e. Dominant spatial patterns of the forced response

The vector $\tilde{\mathbf{e}}_1$ represents an optimal filter for extracting the time evolution of the forced response, but it is not itself an estimate of the leading eigenvector of $\mathbf{C}_F$, so we still do not have an estimate of the spatial characteristics of the leading mode(s) of the forced response. This is given by $\hat{\mathbf{E}} = \mathbf{F}^{(-1)}\mathbf{E}'$, where $\mathbf{F}^{(-1)}$ is the pseudo-inverse of the prewhitening operator. To understand why this is the case suppose, for simplicity, that $\mathbf{C}_F$ is of rank one (i.e., the forced response consists of a single pattern $\mathbf{g}$). We can then write:

$$\mathbf{C}_F = \mathbf{g}\mathbf{g}^{\mathrm{T}} \tag{18}$$

$$\mathcal{E}(\mathbf{C}'_M) = \mathbf{F}^{\mathrm{T}}\mathbf{g}\mathbf{g}^{\mathrm{T}}\mathbf{F} + \mathbf{I} \tag{19}$$

$$= \mathbf{g}'\mathbf{g}'^{\mathrm{T}} + \mathbf{I}. \tag{20}$$

Provided we have sufficient data that $\mathbf{C}'_M \simeq \mathcal{E}(\mathbf{C}'_M)$ then because adding $\mathbf{I}$ to any matrix does not change its eigenvectors,

$$\mathbf{e}'_1 = \alpha\mathbf{g}' \tag{21}$$

$$= \alpha\mathbf{F}^{\mathrm{T}}\mathbf{g} \tag{22}$$

$$\hat{\mathbf{e}}_1 = \alpha\mathbf{F}^{(-1)}\mathbf{F}^{\mathrm{T}}\mathbf{g} \tag{23}$$

$$= \alpha\mathbf{E}_N^{(\kappa)}\mathbf{E}_N^{(\kappa)\mathrm{T}}\mathbf{g}, \tag{24}$$

where α is a normalization factor. The last equality follows from the definition of $\mathbf{F}$ and $\mathbf{F}^{(-1)}$. The vector $\hat{\mathbf{e}}_1$ therefore resembles $\mathbf{g}$ to the extent that $\mathbf{g}$ can be represented by the κ highest-ranked EOFs of the noise.

The vector $\hat{\mathbf{e}}_1$ provides an estimate of the first EOF of $\mathbf{X}_F$, the mean evolution of a hypothetical infinite ensemble, in the case that $\mathbf{X}_F$ consists only of a single pattern. Since we only concern ourselves here with the highest-ranked patterns, we could leave the discussion at this point, but for completeness, we conclude by considering the case where $\mathbf{X}_F$ consists of μ spatial patterns, where $\mu \ll l$. In this case, $\mathbf{C}_F$ will have μ nonzero eigenvalues, and the μ highest-ranked columns of $\hat{\mathbf{E}}$ provide an estimate of the subspace in which the corresponding eigenvectors lie. They will not provide an estimate of the eigenvectors themselves, since $\hat{\mathbf{E}}$ is not an orthonormal basis in the conventional (Euclidean) sense. Specifically,

$$\hat{\mathbf{E}}\hat{\mathbf{E}}^{\mathrm{T}} = n\hat{\mathbf{C}}_N \tag{25}$$

$$n\hat{\mathbf{E}}^{\mathrm{T}}\hat{\mathbf{C}}_N^{(-1)}\hat{\mathbf{E}} = \mathbf{I}. \tag{26}$$

Similar orthogonality conditions apply to $\tilde{\mathbf{E}}$, with $\hat{\mathbf{C}}_N$ replaced by $\hat{\mathbf{C}}_N^{(-1)}$. In the notation of Hasselmann (1997), $\hat{\mathbf{E}}$ and $\tilde{\mathbf{E}}$ represent co- and contravariant tensors under the metric defined by $\hat{\mathbf{C}}_N^{(-1)}$: that is, $\tilde{\mathbf{E}} = \mathbf{F}\mathbf{F}^{\mathrm{T}}\hat{\mathbf{E}} = \hat{\mathbf{C}}_N^{(-1)}\hat{\mathbf{E}}$. Allen and Smith (1997) discuss how an orthonormal basis may be extracted from the leading S/N-maximizing EOFs, but since we will be considering only a single pattern, this need not concern us here.

The contours of $\hat{\mathbf{e}}_1$ (Fig. 3e) can be interpreted physically because they represent the estimate of the dominant eigenvector of $\mathbf{C}_F$. We discuss the implications of these patterns in the next section. Before doing so, we note that a more direct way of obtaining $\hat{\mathbf{E}}$ is by computing the patterns of regression coefficients of the original ensemble mean, filtered by projection onto the first κ EOFs of internal variability, onto the optimized PCs. Using the SVD of the prewhitened ensemble mean $\mathbf{F}^{\mathrm{T}}\mathbf{X}_M = \mathbf{E}'\mathbf{\Lambda}'\mathbf{P}^{\mathrm{T}}$ and the definition of $\tilde{\mathbf{E}}$ one can write

$$\mathbf{\Lambda}'\mathbf{P}^{\mathrm{T}} = \tilde{\mathbf{E}}^{\mathrm{T}}\mathbf{X}_M. \tag{27}$$

Pre- and postmultiplying by $\hat{\mathbf{E}}$ and $\mathbf{P}$ yields

$$\hat{\mathbf{E}}\mathbf{\Lambda}' = \hat{\mathbf{E}}\tilde{\mathbf{E}}^{\mathrm{T}}\mathbf{X}_M\mathbf{P} \tag{28}$$

$$= \mathbf{E}_N^{(\kappa)}\mathbf{\Lambda}_N\mathbf{E}'\mathbf{E}'^{\mathrm{T}}\mathbf{\Lambda}_N^{(-1)}\mathbf{E}_N^{(\kappa)\mathrm{T}}\mathbf{X}_M\mathbf{P} \tag{29}$$

$$= \mathbf{E}_N^{(\kappa)}\mathbf{E}_N^{(\kappa)\mathrm{T}}\mathbf{X}_M\mathbf{P}, \tag{30}$$

which since $\mathbf{P}$ is orthonormal, is equivalent to regressing the filtered mean, $\mathbf{E}_N^{(\kappa)}\mathbf{E}_N^{(\kappa)\mathrm{T}}\mathbf{X}_M$, onto the optimized PCs. We will subsequently be discussing patterns of regression coefficients of various variables onto these optimized PCs, so it is helpful to bear this link in mind.

4. Atmospheric response to North Atlantic SST changes

a. Removing remote effects of ENSO

The principal origin of the dominant signal identified in the previous section is the remote response to El Niño. The characteristic time series (Fig. 3b) contain strong interannual fluctuations and the optimal filter patterns (Fig. 3a), which extract these from the ensemble mean MSLP are most pronounced over Central America where the atmosphere is strongly influenced by the Southern Oscillation. If we map the regression coefficients of SST anomalies on the time series obtained in the previous section (Fig. 4a, contour lines) and the fraction of total SST variance explained by the regression (shaded), the tropical Pacific origin of the signal is clear.

In this study the focus of our interest is the atmospheric response to changes in Atlantic Ocean SSTs, so we reduced the component of atmospheric variability that is due to the remote impact of tropical Pacific SST anomalies (especially those that are associated with ENSO) by individually computing the regression coefficients at each grid point of the seasonal mean MSLP, wind stress, and heat flux anomalies from every ensemble member with the Niño-3 SST index (5°N–5°S, 150°–90°W) and retaining only the residual part of the data for our analysis. Repeating the signal-to-noise maximizing analysis of section 3 with the ''ENSO removed'' MSLP data, the characteristic time series (Fig. 5b) change considerably. They contain much less interannual and more pronounced decadal fluctuations than before. As can be inferred from Fig. 4b, the detected atmospheric response can now mainly be attributed to Atlantic Ocean SST anomalies, which exhibit the tripole structure reported in various other studies of model data and observations (e.g., Deser and Blackmon 1993; Grötzner et al. 1998; Sutton and Allen 1997). The remote effect of ENSO-related SST anomalies in the tropical Pacific Ocean was successfully removed. In so doing a portion of the common signal was also removed from the data, causing a slight decrease in the signal-to-noise ratios (Fig. 5c). Figure 5d, however, shows that the variance of $\mathbf{X}_M$ in the direction of $\tilde{\mathbf{e}}_1$ still is almost one order of magnitude larger than the corresponding variances of the $\mathbf{X}_{Nk}$. The estimate of the dominant MSLP response, $\hat{\mathbf{e}}_1$, (Fig. 5e) now has a dipole-structure similar to that shown in Fig. 1. We now know, however, that different parts of this pattern are associated with very different ratios of signal-to-noise, indicating for different regions different chances to successfully predict long-term MSLP changes from prescribed SSTs. The most predictable variability is found over the tropical and subtropical Atlantic, as shown by the optimal filter pattern, $\tilde{\mathbf{e}}_1$. The temporal evolution of $\tilde{\mathbf{e}}_1$ is, in both winter and spring, detectable at the 95% level. In spring, in contrast to the estimated forced response from standard PCA, the evolution of this pattern is even consistent between the individual ensemble members, in the sense defined above (Table 1, lines 7 and 8).

The main result to stress here is that we have extracted a detectable, and in spring even consistent, response to SST and sea-ice anomalies in the modeled North Atlantic MSLP data. Regression analysis suggests this sig-

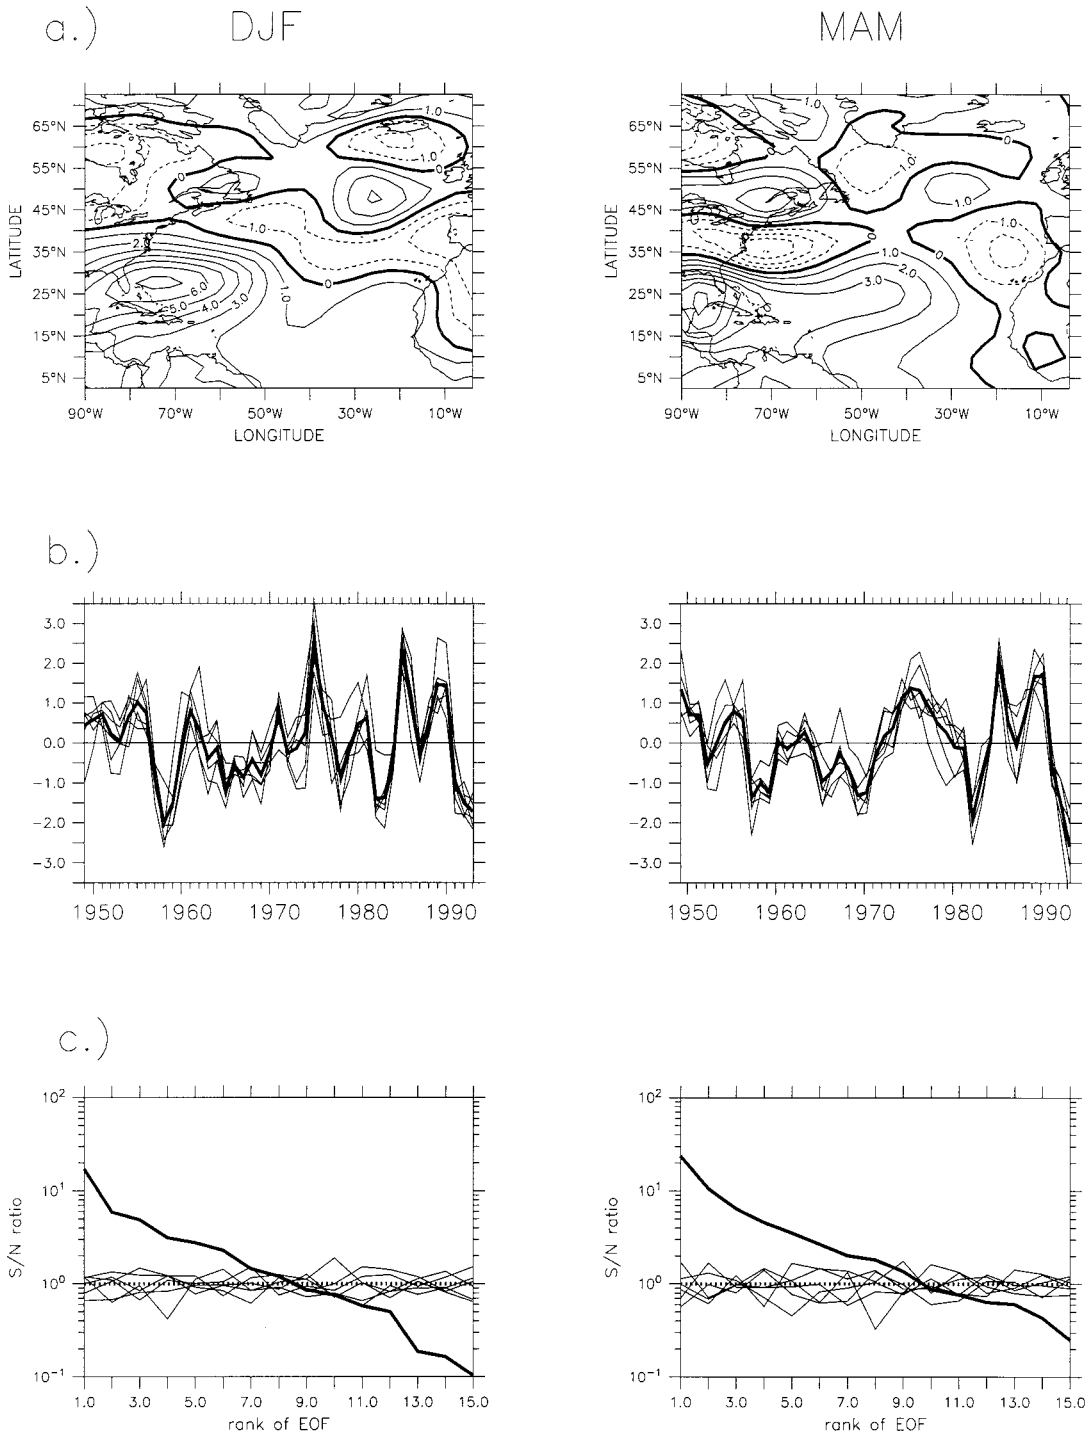

FIG. 3. (a) Normalized optimal filter patterns, $\tilde{\mathbf{e}}_1$, for winter and spring MSLP. (b) Projections of ensemble mean (heavy line) and individual ensemble members (thin lines) MSLP onto the patterns shown in (a). (c) Signal-to-noise ratios of leading 15 signal-to-noise maximizing EOFs (heavy line) and corresponding values Λ'^2_{Nk} for the individual noise realizations $\mathbf{X}_{Nk}$ (thin lines). (d) Variance of ensemble mean (heavy line) and individual noise realizations (thin lines) in the directions

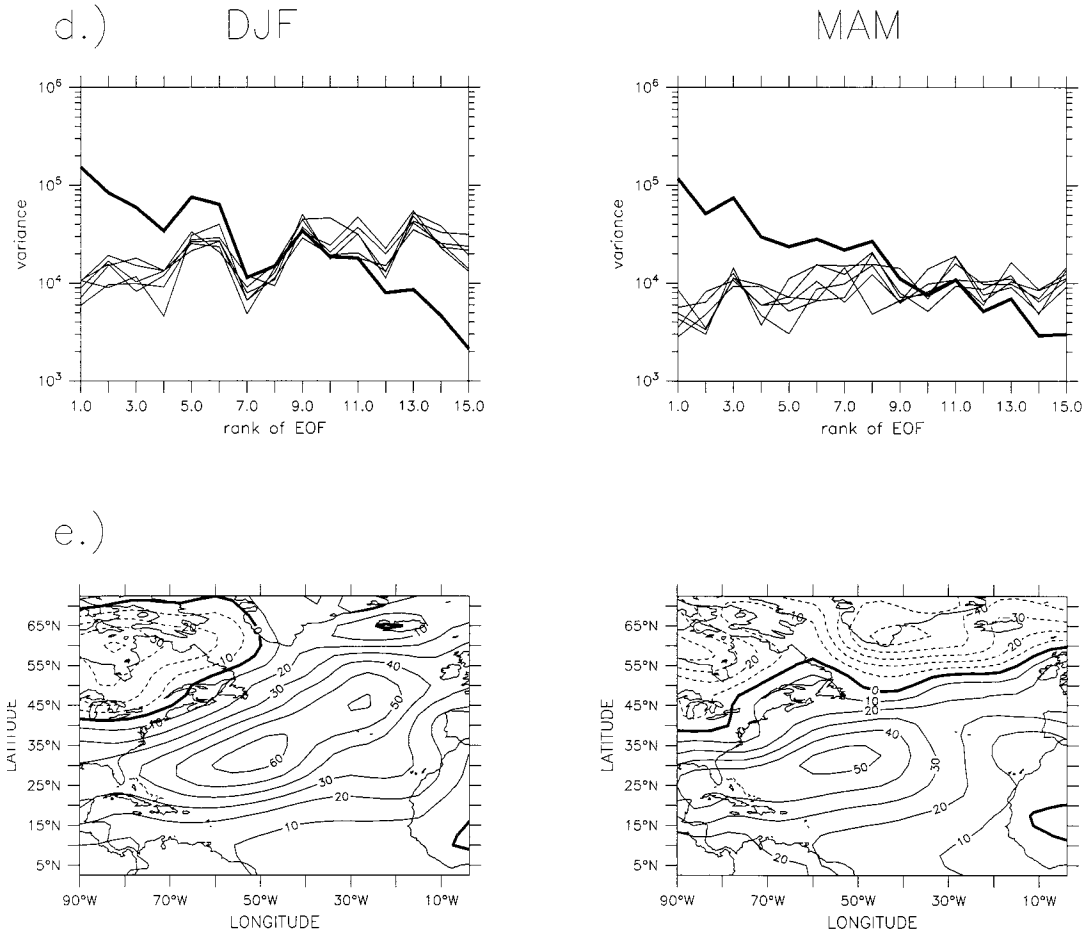

FIG. 3. (*continued*) defined by $\tilde{\mathbf{E}}$. (e) Signal-to-noise maximizing dominant eigenvector, $\hat{\mathbf{e}}_1$, of $\mathbf{C}_F$, estimating the true forced response. The time series are standardized to unit variance and the patterns are multiplied by the corresponding singular value so that the pairs of patterns and time series retain the correct variance.

nal arises as a response to fluctuations in a tripole pattern of North Atlantic SST. We will discuss the form of the detected response and its implications for coupled mechanisms of decadal climate variability in detail in section 4d. Before that, we discuss how the identified model response relates to observations and address the stability of the statistical method used.

b. Comparison with observations

Does the MSLP response we have detected in our model bear any resemblance to fluctuations in MSLP seen in the real world? MSLP observations for the period considered have been compiled in the COADS dataset (da Silva et al. 1994). Figure 6a shows the time series obtained when we project these observations onto the first EOFs of the ensemble mean from standard PCA, and the time series obtained when we project the ensemble mean and the ensemble members onto the same EOFs. There is little agreement between the time series derived from the observations and that derived from the ensemble mean. Projections onto the optimal filter patterns, however, yield time series that appear well correlated, both in the case of full MSLP fields (Fig. 6b) and in the ''ENSO removed'' case (Fig. 6c). These results show that not only is there a detectable and consistent response of the atmosphere model to oceanic forcing but that this response is fairly well matched in the observations.

The comparisons shown in Fig. 6 are influenced by both the temporal and the spatial characteristics of our model's response. As a consequence of small systematic errors in the model mean state it is possible that the model responds to the same SST fluctuation as the real atmosphere (so that the temporal evolution of the response is correct), but produces a spatial pattern of MSLP, which is somewhat in error. Such errors may be a cause of some of the discrepancies between the projections of the observed data and the model data onto the optimal filter patterns. Figure 7 shows a comparison between the model and observations that is not influenced by such errors. Here we compare the time series

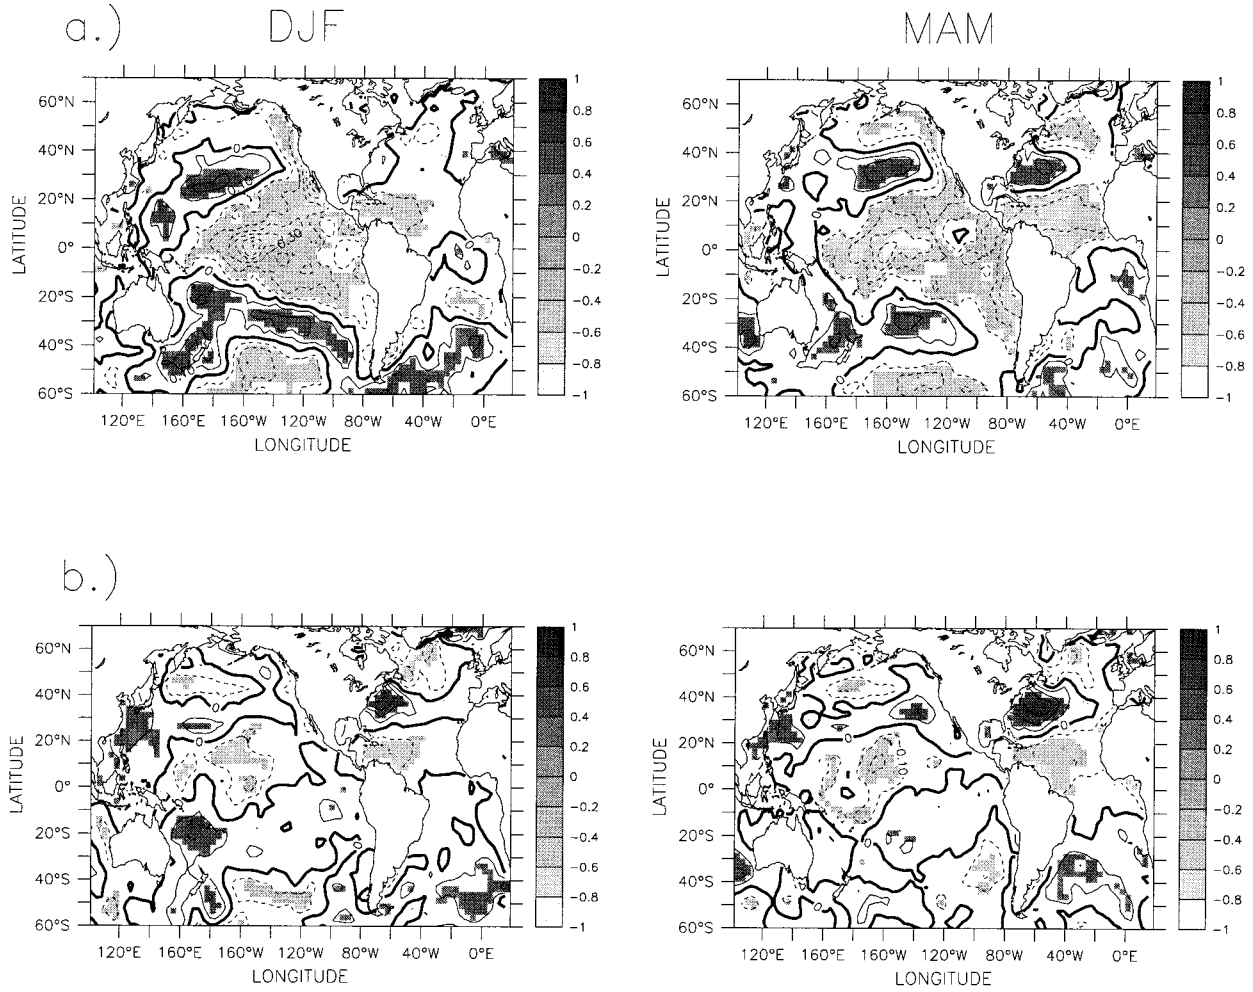

FIG. 4. (a) Regression coefficients (K; contour lines) and fraction of total variance explained (shaded contours with sign appropriate to the sign of the regression coefficient, plotted only where significant at the 95% confidence level) by the regression of SST anomalies onto time series shown in Fig. 3b (heavy line). (b) Same as (a) but for the regression onto time series shown in Fig. 5b.

associated with the most predictable mode of HadAM1 MSLP variability ($\mathbf{p}_1$ obtained from the optimal filter) with the index time series of *observed* SST and MSLP variations that were found by Sutton and Allen (1997) to be characteristic of the propagation of temperature anomalies across the North Atlantic. We stress that no observational data (other than, indirectly, the SSTs used to forced the model) have been used to derive the HadAM1 MSLP optimized time series (in contrast, e.g., to a canonical correlation analysis). The agreement (significant at the 95% confidence level) shows that the common pattern obtained through the signal-to-noise maximizing analysis of model MSLP exhibits a temporal evolution that is characteristic of the underlying ocean dynamics. It suggests that the observed correlations between SST and MSLP fluctuations must reflect, at least in part, an influence of the ocean on the atmosphere. The nature of this influence will be discussed further in section 4d.

c. *Sensitivity to truncation and variable used in optimization*

Since the method we used to specify a truncation level was only heuristic, we must next address the sensitivity of the detected atmospheric response to the truncation κ in the definition of $\mathbf{F}$ and $\mathbf{F}^{(-1)}$ [the same problem arises when optimal fingerprints are used for the detection of anthropogenic climate change (see, e.g., Hegerl et al. 1996)]. The characteristic temporal evolution that results from the signal-to-noise maximizing analysis shows only marginal sensitivity to variations in the truncation level κ as can be seen in Fig. 8a. This is reassuring because, as derived in section 3, we estimate the spatial structure of the dominant forced response by computing maps of the regression coefficients of the different atmospheric fields regressed onto this time series. Our key results, discussed in detail in the next section, are therefore stable against the choice of truncation.

So far the analysis has been applied only to MSLP data. We also applied the whole signal-to-noise maximizing algorithm to other atmospheric variables, including wind stress and surface heat fluxes, and always detected very similar characteristic time series (Fig. 8b). The agreement between these series is remarkable and supports the hypothesis that the algorithm has identified a physically coherent signal. For consistency we will henceforth always refer to the time series shown in Fig. 5b (thick lines) as the ''characteristic temporal evolution'' of the forced response.

d. Spatial characteristics of the response

We have established that our ensemble of atmosphere model integrations exhibits common decadal fluctuations in the subtropical North Atlantic region. The temporal evolution of the common signal resembles a time series that has been related to the slow propagation of temperature anomalies across the North Atlantic. The implication is that dynamical processes in the North Atlantic ocean may be influencing decadal fluctuations in the atmosphere above. We need to ask next how the atmospheric fluctuations might feed back to affect the ocean, and to consider in particular whether our results support the possibility of a coupled ocean–atmosphere mode of decadal variability. To address these questions we examine the spatial characteristics of the response in several different variables.

Figure 9 shows maps of the regression coefficients (contour lines) of anomalous SSTs and ensemble means of: MSLP, wind stress, net surface heat flux, and wind stress curl onto the characteristic temporal evolution of the forced response (Fig. 5b). The shading indicates the fraction of total ensemble mean variance explained by the regression. Corresponding regression maps were also computed for the individual ensemble members. Shaded areas indicate regions where regression coefficients of the ensemble mean were found to be significantly different from zero (i.e., ''detectable'') at the 95% confidence level based on a Student's t-test. In line with the findings of Grötzner et al. (1998), MSLP fluctuations associated with the detected response are dominated by a large-scale dipole structure (Fig. 5e and Fig. 9b). In DJF the dipole is located in the northeastern part of the Atlantic basin, while in MAM it is shifted southwestward. The fraction of total MSLP variance explained by the regression is largest over the tropical and subtropical North Atlantic, in line with our expectations from the optimal filter pattern (Fig. 5a) and also with the observational analysis of Sutton and Allen (1997). The higher level of internal variability at high latitudes means that the atmospheric response is less strong here and in many regions is not detectably different from zero. We note that gridpoint-based analysis of variance, ANOVA (Rowell et al. 1995; Davies et al. 1997; Rowell 1998), yields the same regions of significantly nonzero MSLP response (not shown). An estimate of the spatial characteristics and temporal evolution of the response, however, cannot be inferred from this gridpoint statistic.

The spatial patterns that dominate the forced response (Fig. 9b) closely resemble the spatial patterns that dominate the internal variability (not shown). This result suggests that the characteristics of the internal variability strongly influence the response of the system to perturbations (specifically changes in SST). Such behavior is typical of systems that exhibit regime-like behavior (e.g., Lorenz 1963; Robertson et al. 1999, manuscript submitted to *J. Atmos. Sci.*; Renshaw et al. 1998). We have not carried out an analysis of weather regimes in our data, but a possible explanation for this result is that the leading mode of internal variability is associated with transitions between two regimes, and that the SST influence the relative amount of time the system spends in these two regimes. This said, the fact that the optimal filter pattern differs from the leading pattern in the forced response and the leading pattern in the internal variability tells us that the influence of SST is not limited to this effect alone.

The wind stress response (Fig. 9c) shows the main features one would anticipate from the MSLP response assuming geostrophic balance. In spring, when the response is most consistent, there is a tendency for intensified westerlies (around 55°N) and, more consistently, for intensified trades (around 15°N) in association with anomalously warm SSTs off the U.S. coast and anomously cold SSTs in the tropical–subtropical Atlantic, and in the North Atlantic south of Greenland (Fig. 9a). Through their effects on surface heat fluxes and Ekman fluxes we would expect these wind anomalies to affect local SSTs. We show the heat fluxes (positive into the ocean) associated with the common response in Fig. 9d. In considering these, however, we need to keep in mind that in our experiments the SST is prescribed. This artificial constraint could lead to an overestimate of heat flux anomalies (see Barsugli and Battisti 1998; Saravanan and McWilliams 1998). It is also important to note that we cannot establish from our experiments whether the atmospheric response is induced by the whole SST tripole pattern or only by some part of it. If the latter, then correlated SST fluctuations in one part of the tripole could, in the real world, be forced by the atmospheric response to SST fluctuations in the other parts.

The enhanced westerlies around 55°N in spring will act to cool the local ocean surface through enhanced extraction of heat by turbulent fluxes associated with enhanced advection of cold air from the Labrador Sea and Canada and enhanced wind speeds, and through advection of the mean temperature field by anomalous Ekman currents. Even though the SSTs in this region are already (i.e., are prescribed to be) anomalously cold, Fig. 9d shows that there is nonetheless an anomalous heat flux out of the ocean in MAM. A breakdown of the surface heat flux into its individual components (not

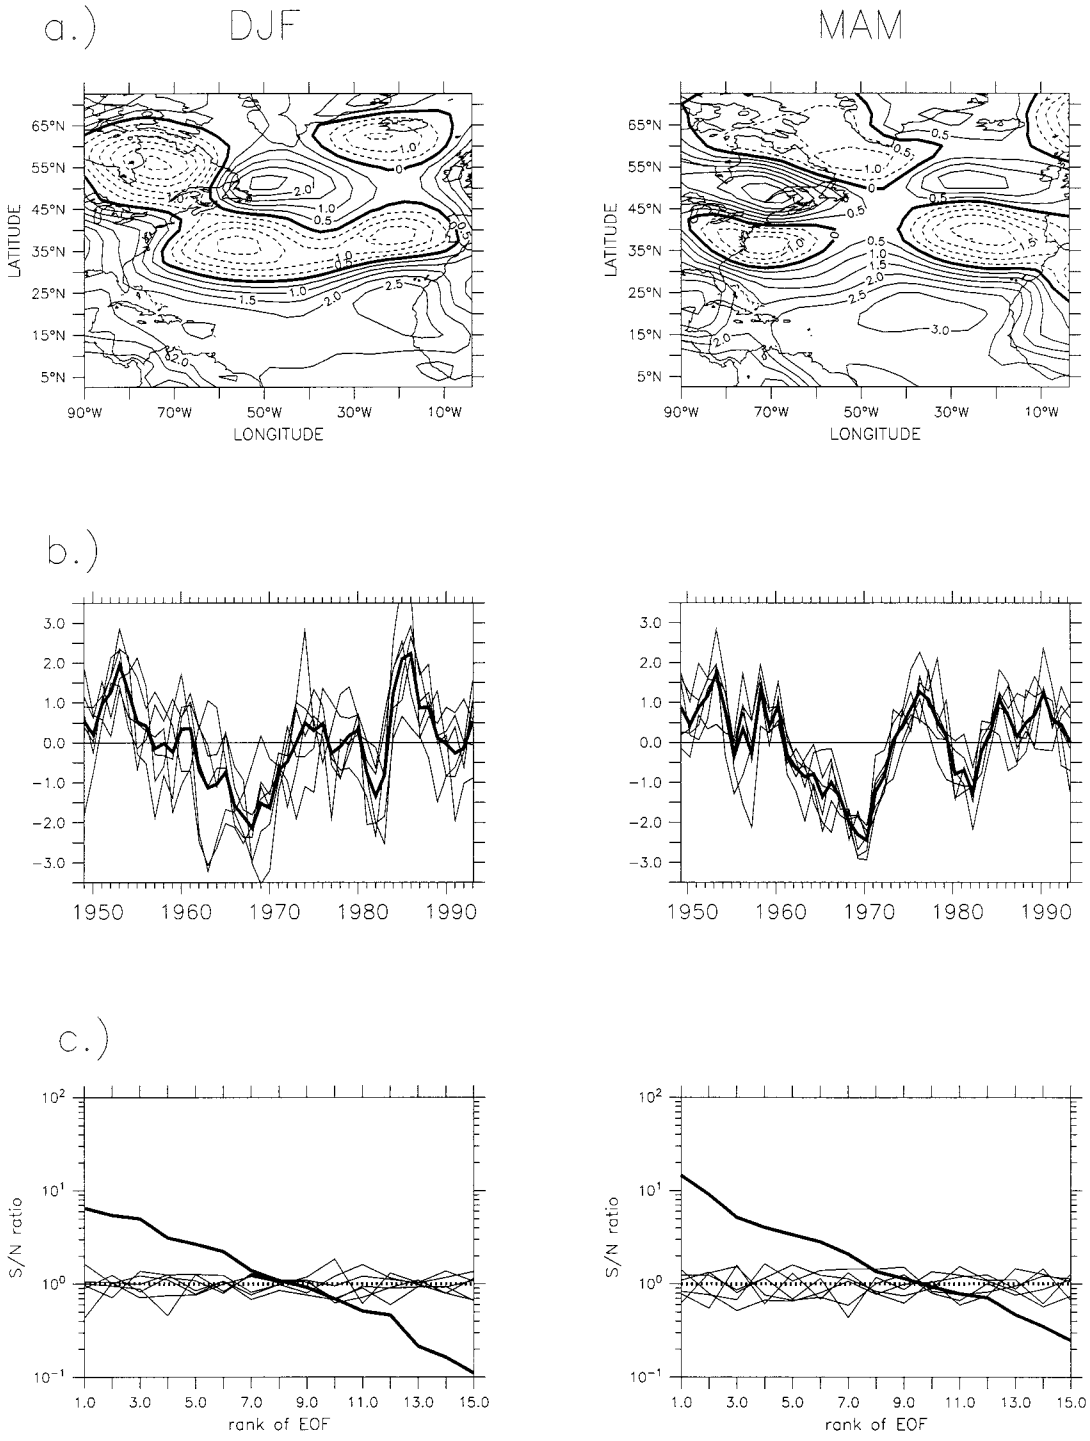

FIG. 5. Same as Fig. 3 but for ''ENSO removed'' MSLP data.

shown) shows that the major contributions are from sensible and latent heat fluxes, as we anticipated.

If this northern component of the SST tripole is playing an active role in the forcing, then the atmospheric response in this region seems to indicate a weak positive feedback, as noted by Grötzner et al.(1998). Based on the t-test, the ensemble-mean response in the shaded area is readily detectable (significantly different from zero at the 95% confidence level). The much more stringent test of consistency indicates that, if we forced the

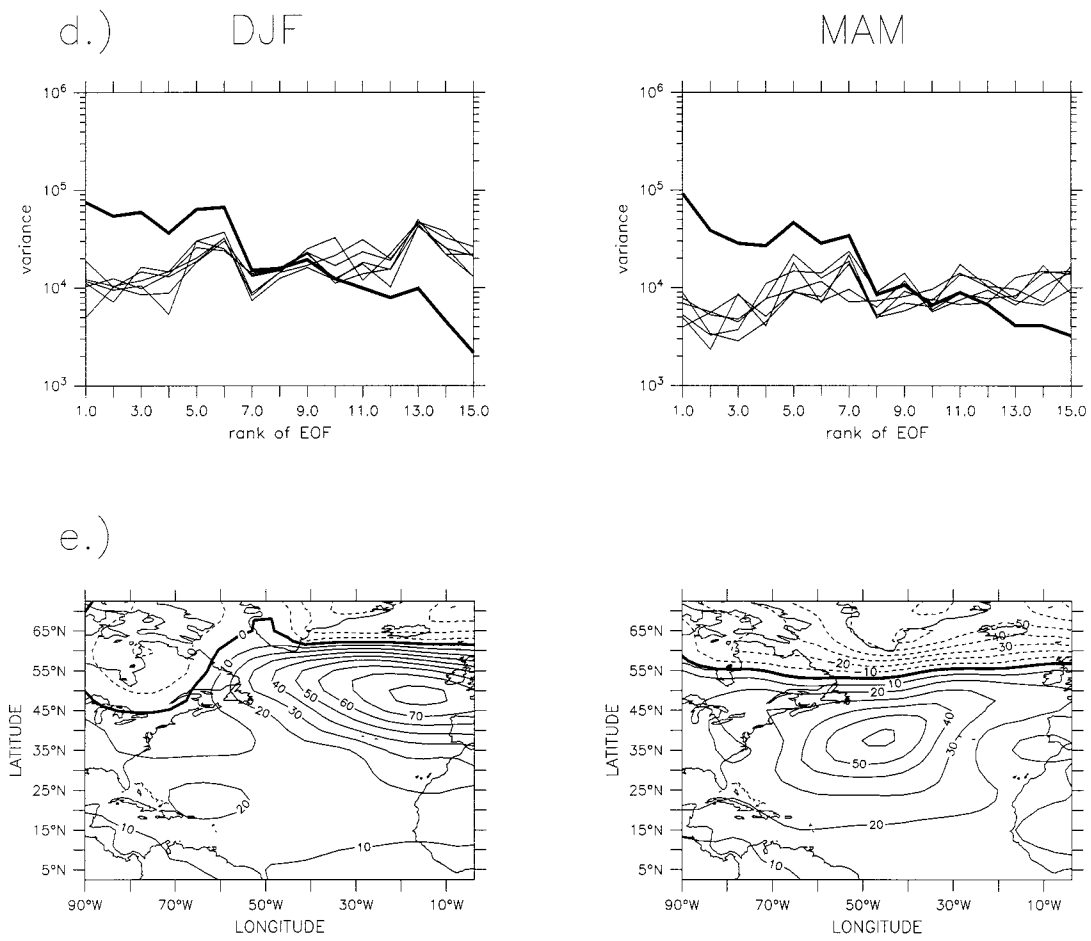

FIG. 5. (*Continued*)

atmosphere model again with the same SSTs, we would obtain zonal wind stress (net surface heat flux) anomalies in this region positively (negatively) correlated with the characteristic temporal evolution shown in Fig. 5b in approximately 90% of cases. The fact that the response would be of the wrong sign in over 10% of cases even after filtering by projection onto the characteristic time series indicates that any feedback cannot be particularly robust.

If, on the other hand, the atmospheric response is forced entirely by the more southerly components of the SST tripole, then these results should not be interpreted as evidence of an atmospheric feedback. Rather, they would suggest that forced fluctuations in high-latitude North Atlantic SST could arise as a remote response to lower-latitude SST changes.

In phase with the enhanced westerlies around 55°N, enhanced trade winds are found over the tropical–subtropical ocean, again especially in spring (Fig. 9c). In contrast to the westerly belt, in this region even the sign of the zonal and meridional wind stress regression coefficients is found in spring to be consistent at the 95% confidence level. One would expect that stronger trades bring cooler air from higher latitudes and enhance the turbulent heat fluxes, leading to anomalous cooling of the upper tropical–subtropical ocean. Even when considering that the rate of the upper-ocean cooling is inversely proportional to the depth of the mixed layer (which is shallower in the Tropics than in midlatitudes), there is little evidence in MAM of surface heat loss in this region (Fig. 9d). This is presumably because the anomalously cold SSTs offset the effect on the heat fluxes of the enhanced trades. In the trade belt, therefore, the heat flux fields suggest that the atmospheric response does not generate a positive feedback on SST anomalies. Neither, however, is there clear indication in MAM of a negative feedback. Without the changes in atmospheric circulation (i.e., the changes in the strength of the trades) a negative feedback would be expected (as is found in DJF). Therefore it appears that the atmospheric circulation response has served to decrease, or even eliminate, the damping of tropical Atlantic SST anomalies by anomalous surface heat fluxes. This effect should, inturn, increase the persistence of SST anom-

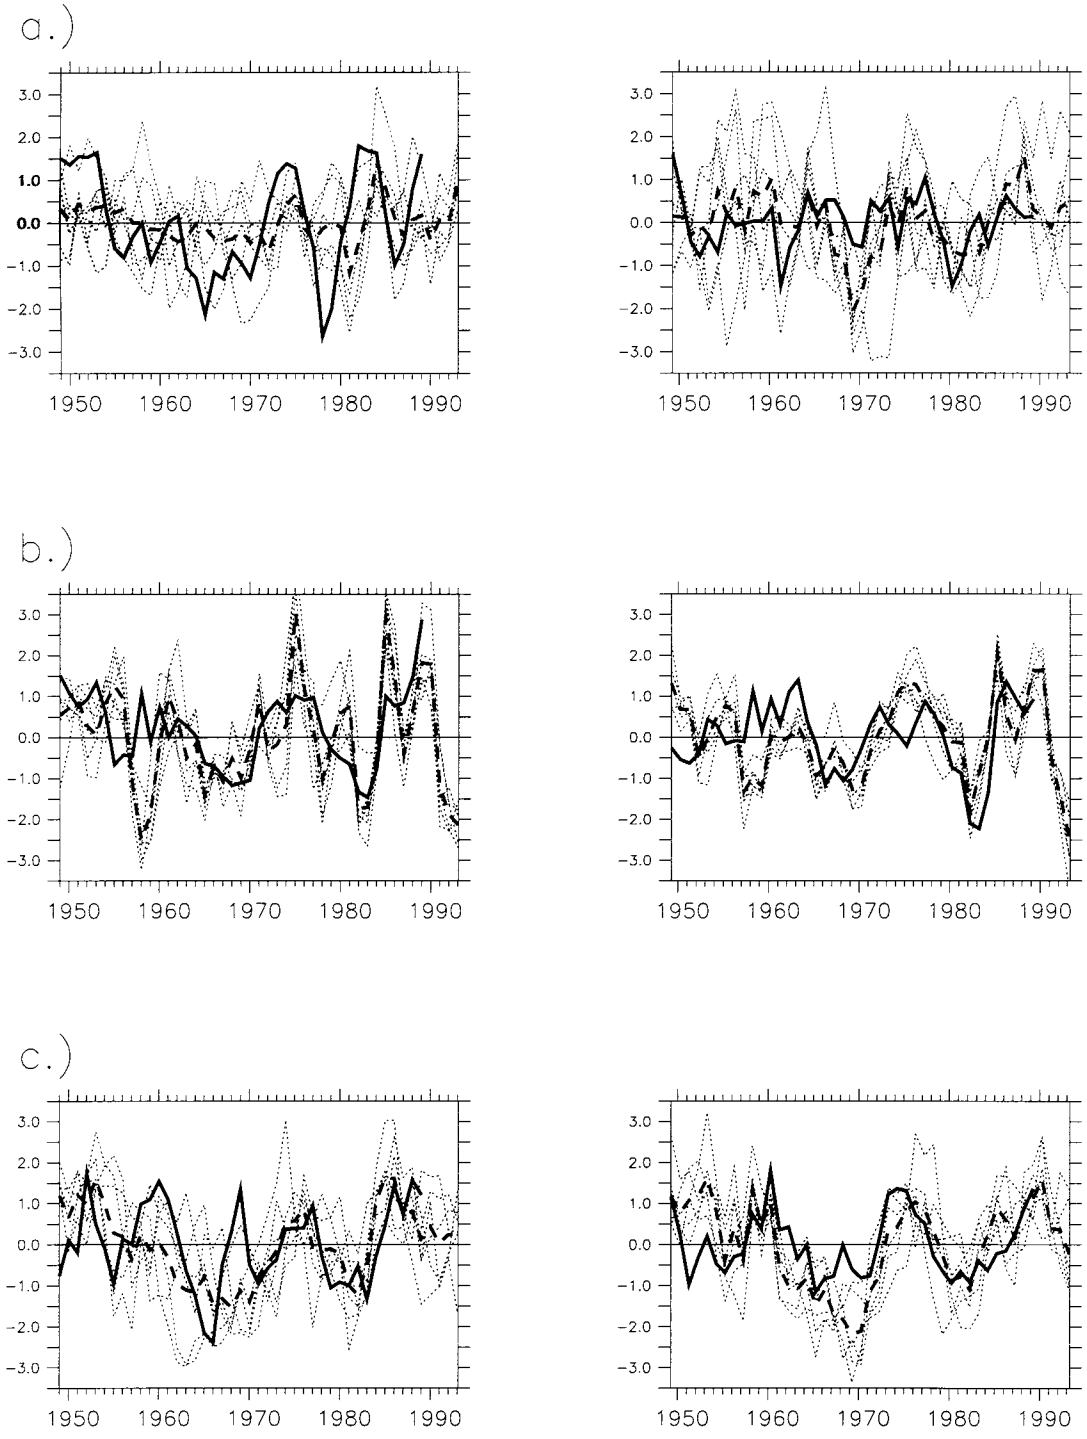

FIG. 6. Projections of winter (left panels) and spring (right panels) mean sea level pressure from individual ensemble members (thin dotted lines), the ensemble mean (dashed line), and observations (COADS—solid line) onto (a) patterns show in Fig. 1a, (b) patterns shown in Fig. 3a, and (c) after ''ENSO removing'' onto patterns shown in Fig. 5a. As the COADS dataset contains only data over sea, all projections were computed over sea points only.

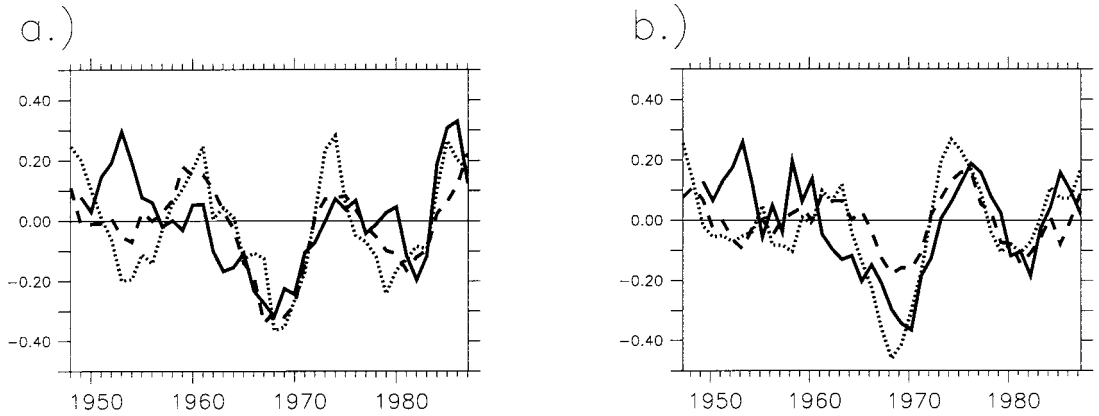

FIG. 7. Characteristic time series of dominant MSLP response (solid line) in (a) winter and (b) spring together with index time series (averaged over 15°–25°N; 40°–60°W) of low-pass filtered observed MSLP anomalies (dashed line) and index time series (averaged over 25°–35°N; 70°–80°W) of observed SST anomalies (dotted line). The magnitude of the curves has been scaled. The correlation coefficients between the detected signal and the observed SST as well as the detected signal and the observed MSLP in spring are 0.48 and 0.62, respectively. In winter the corresponding correlation coefficients are 0.61 and 0.51. These correlations are all significant within the 95% confidence limit.

alies in this region compared to their persistence in a scenario where there is no atmospheric circulation response. Finally, we note that, in addition to their influence on surface heat fluxes, the anomalous trades could influence tropical Atlantic SSTs through other mechanisms. A positive feedback could arise through the tendency of the enhanced trades to cause enhanced mixing, and hence cooling, of the ocean mixed layer. A negative feedback could arise through the advection of the mean SST field by anomalous Ekman currents.

In the final lobe of the SST tripole, the western Atlantic around 30°–40°N, the wind fluctuations in the atmospheric response are not reliably detectable either in DJF or MAM. In line with the results of Grötzner et al. (1998), the heat flux fields suggest a negative feedback. In winter, advection of cold air from the North American continent over anomalously warm water leads to enhanced latent and sensible heat loss.

An interesting feature in spring is found at around 50°W, 30°N. Here, surface heat flux anomalies arise that are both detectable and consistent. These anomalies are mainly due to reduced latent heat loss and enhanced solar radiation. The location and sign suggests that these heat flux anomalies might help to propagate SST anomalies from the Gulf Stream region eastward.

We summarize the local feedbacks of the atmospheric response onto the SST tripole as follows: we have detected a possible weak positive feedback in the North Atlantic south of Greenland and a weak negative feedback in the western Atlantic off the U.S. coast. In the tropical Atlantic we found a negative feedback in DJF, but in MAM the anomalous heat fluxes are close to zero.

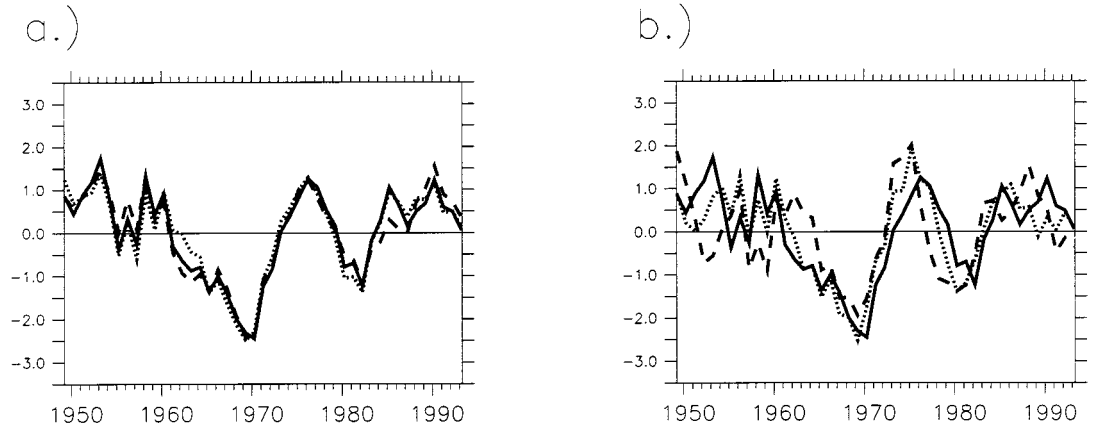

FIG. 8. (a) First principal components of ''ENSO removed'' spring MAM MSLP obtained from the signal-to-noise maximizing EOF analysis for truncation at $\kappa = 14, 19, 24$ noise EOFs (dotted, solid, dashed line, respectively). (b) First principal components of ''ENSO removed'' spring MSLP (solid lines), zonal wind stress (dotted lines), and meridional wind stress (dashed lines) obtained from the signal-to-noise maximizing EOF analysis.

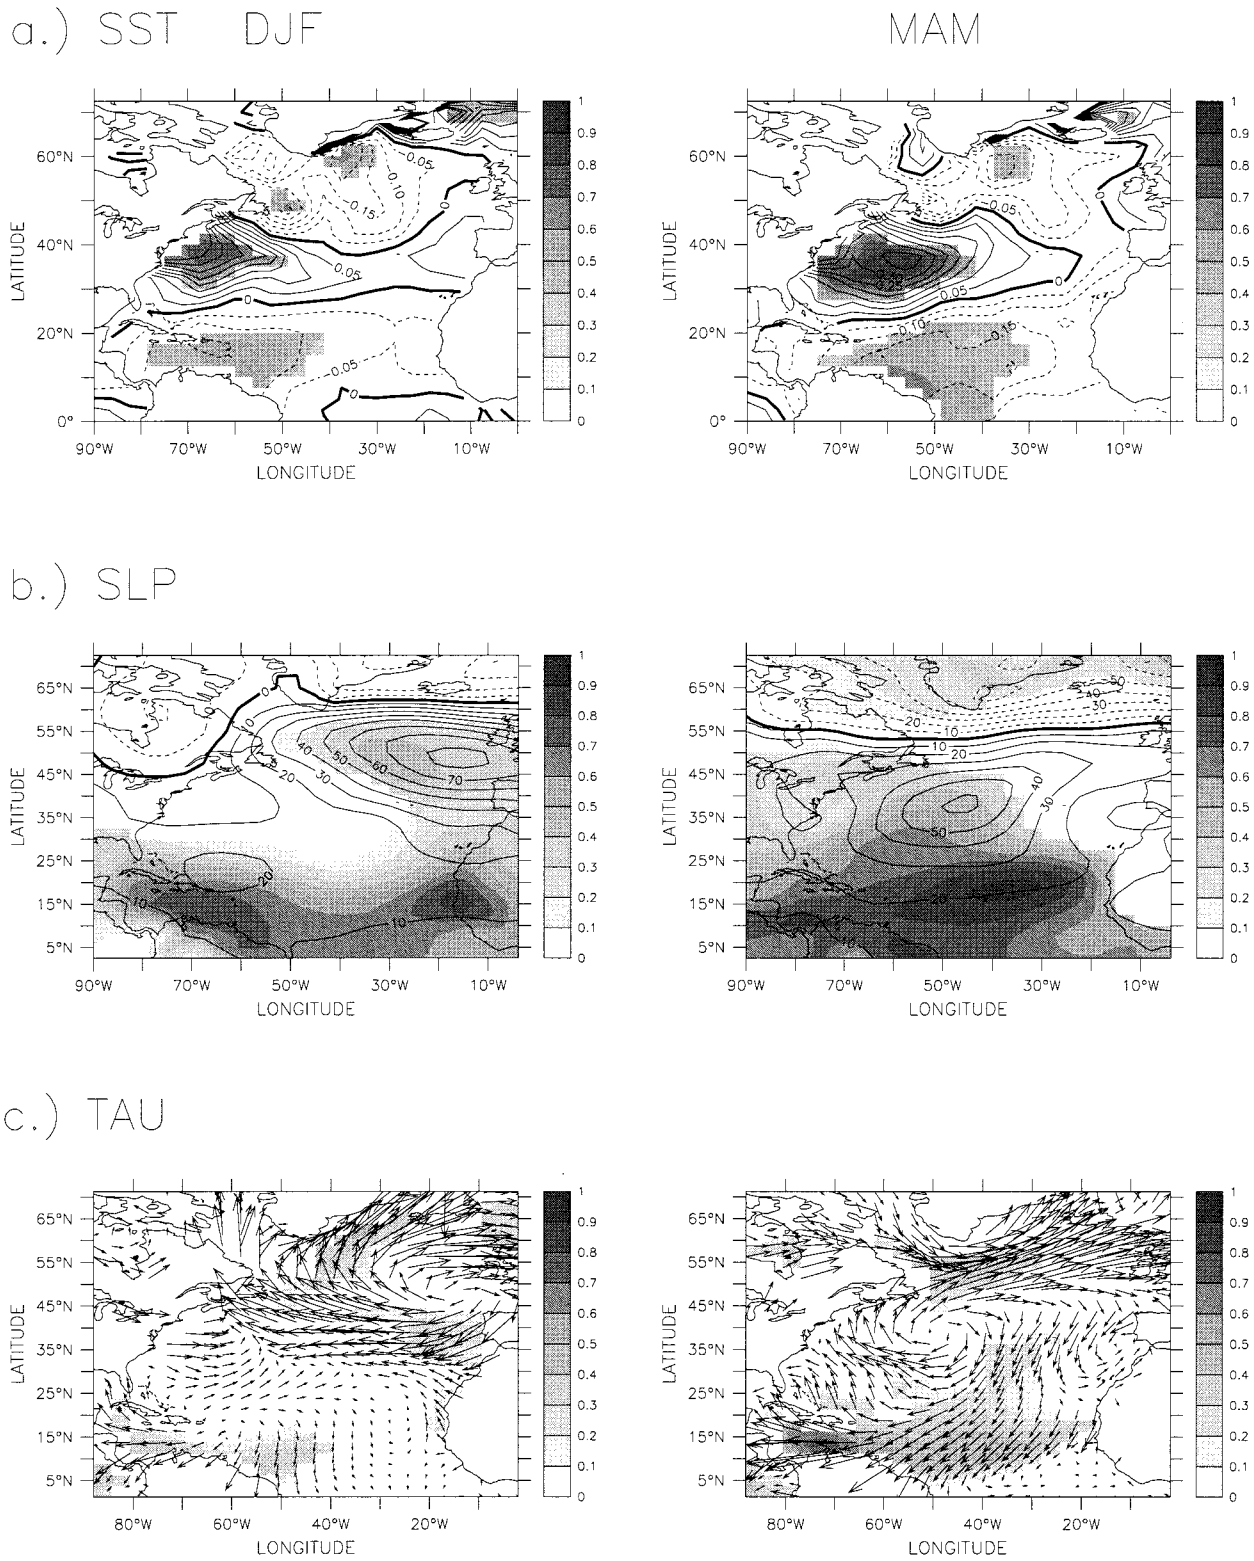

FIG. 9. Regression coefficients (contour lines) and fraction of total variance explained by (shaded contours, plotted only where significant at the 95% confidence level) regression of observed SST anomalies (K) onto time series shown in Fig. 5b (heavy line). Regression coeffi-

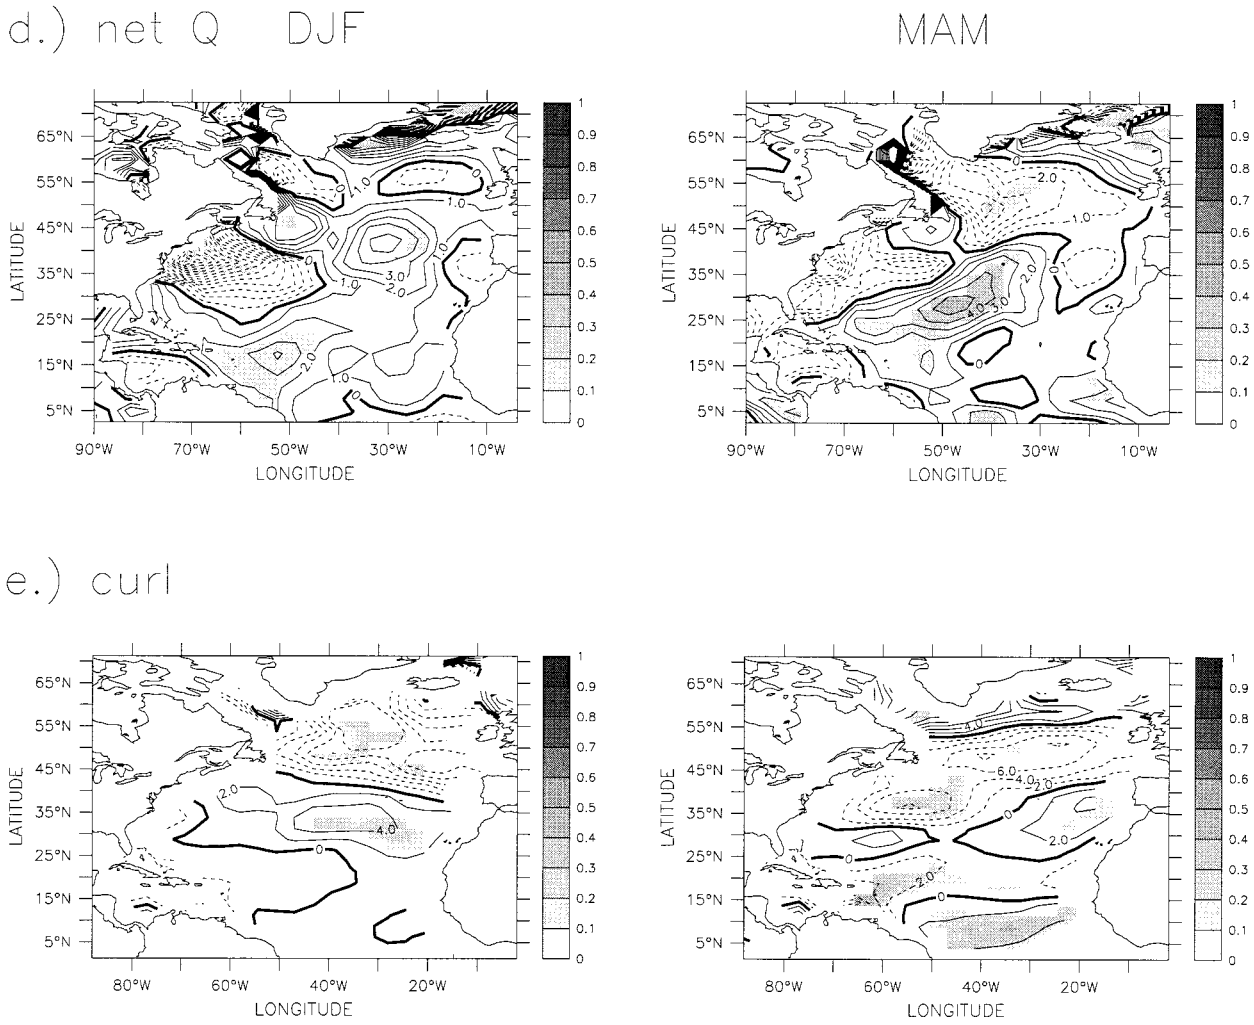

FIG. 9. (*Continued*) significant at the 95% confidence level) of ''ENSO removed'' (b) mean sea level pressure (Pa), (c) wind stress (Pa), (d) net surface heat flux (Wm^2) and (e) wind stress curl (10^{-8} Pa m^{-1}) onto time series shown in Fig. 5b (heavy line).

The latter suggests that the atmospheric response acts to enhance the persistence of SST anomalies in this region. In all three regions the consistency of the feedbacks is marginal at best. This tends to suggest that if these feedbacks play a role in any coupled mode of variability then such a mode would not be robust against disruption by internal atmospheric variability. We also note again that the very concept of a feedback requires assumptions about exactly which regions in the SST distribution are forcing the atmospheric response.

As well as providing local instantaneous feedbacks, it has also been proposed (Latif and Barnett 1994; Grötzner et al. 1998) that the atmospheric response could influence the ocean in such a way as to give a delayed nonlocal feedback onto SST and a prefered timescale for a coupled ocean–atmosphere mode of variability. In these theories it is argued that the atmosphere responds especially to SST anomalies in the Gulf Stream region. If there is a phase reversal of SSTs here, this could be communicated to other regions of the North Atlantic through the atmospheric response. The response that we detected in spring, if it is primarily a response to SST anomalies in the Gulf Stream region, does have features that could provide such communication: the anomalous westerlies and trades could force a phase reversal of SST anomalies in the North Atlantic and in the tropical Atlantic, respectively.

Grötzner et al. (1998) describe two possible mechanisms for the delayed feedback to Gulf Stream SSTs. In the first, temperature anomalies in the tropical ocean are advected westward along the southern branch of the subtropical gyre and feed into the western boundary current. There they replace anomalies of the opposite sign, giving a phase reversal. Because it relies on advection by the mean ocean circulation, our results cannot be used to test this theory further. The second hypothesis, by contrast, can be considered further. It requires the atmospheric response to include anomalies in wind

stress curl. After a lag associated with the propagation of baroclinic Rossby waves, these anomalies force variations in the strength of the subtropical gyre circulation, which modulate SST especially in the Gulf Stream region. The regression pattern of the ensemble mean wind stress curl (Fig. 9e) is characterized by small-scale structures. To generate Rossby waves that spin down the subtropical gyre circulation (giving a negative feedback), positive wind stress curl anomalies are required south of about 35°N. There is a little evidence for anomalies of the right sign in Fig. 9e but these are of small spatial extent, in general are not detectable, and are nowhere consistent. Our results therefore do not support the hypothesis of a delayed negative feedback supplied by changes in wind stress curl.

5. Summary and conclusions

We have presented an algorithm to estimate the forced response of the atmosphere to observed SST variations using an ensemble of integrations of an atmospheric general circulation model. Forcing the atmospheric model with observed SSTs rather than idealized SST anomalies allows direct comparison with observations but leads to problems of signal-to-noise in a small ensemble. We define the forced response to be estimated as the hypothetical mean evolution of an infinite ensemble of integrations, while recognizing that there are responses (e.g., SST-induced changes in storm-track variability) that would not be captured by this definition.

At the signal-to-noise levels typically encountered in the extratropics, conventional principal component (or EOF) analysis of the mean of a variable such as MSLP from a relatively small (six-member) ensemble yields only an estimate of the true forced response if the internal variability is uncorrelated in space (i.e., is spatially white). This is never the case for atmospheric data. Using a ''prewhitening'' transformation that includes information about the structure of the noise into the EOF analysis, an estimate of the true forced response can be determined.

Our procedure may be summarized thus: (i) we project the ensemble mean onto the EOFs of within-ensemble variability (which provides an estimate of the noise); (ii) we weight the projections to give the same expected noise variance in each; (iii) we perform a standard PCA/EOF analysis of the weighted (prewhitened) data; and (iv) we transform back to ordinary coordinates to interpret the results.

We applied this method to North Atlantic low-pass-filtered seasonal mean MSLP anomalies and detected a common signal. A portion of the common signal was attributed to oceanic changes associated with the El Niño–Southern Oscillation. After suppressing this influence, a detectable response remained in both winter (DJF) and spring (MAM) seasons, with the spring response being not only detectable in the ensemble mean, but also strong enough that the individual ensemble members track the ensemble mean consistently throughout the period considered. The characteristic temporal evolution of this response is also obtained from applying the same technique to other atmospheric variables and from a projection of observed MSLP onto the ''optimal filter'' spatial pattern estimated from the AGCM—indicating it is a genuine signal and not an artefact of the analysis technique or AGCM. This temporal evolution is also well correlated with fluctuations in a tripole pattern of SST anomalies in the North Atlantic, which Sutton and Allen (1997) have identified as one phase of a decadal fluctuation in which propagation of SST anomalies plays a role.

Our best estimate of the spatial pattern associated with the dominant mode of the forced response (which, because of the need for prewhitening, is not the same as the optimal filter) is dominated by a dipole pattern in MSLP. This dipole resembles the NAO more in MAM than in DJF. Only over the tropical–subtropical Atlantic, however, does the forced response account for a substantial fraction of the total variance. Related to the MSLP response there are fluctuations in wind stress, particularly in the westerly belt around 55°N and in the trades. The fluctuations in the trades are more consistent between ensemble members than the fluctuations in the westerlies.

We considered how the atmospheric response might feed back to affect the ocean, in particular the SST tripole. The only region in which we found any evidence of a positive heat flux feedback was the North Atlantic south of Greenland. In the western Atlantic off the U.S. coast the heat flux feedback is negative. In the tropical Atlantic we found a negative heat flux feedback in DJF, and a feedback close to zero in MAM. If the response we detect is forced principally by SSTs in the Gulf Stream region then the form of this response suggests that a change in the sign of Gulf Stream SST anomalies could be communicated in spring to both the North Atlantic and the tropical Atlantic through the atmosphere. This result offers some support for the idea of Grötzner et al. (1998) that a delayed negative feedback to Gulf Stream SSTs, perhaps providing a timescale for a coupled ocean–atmosphere mode of variability, could arise from oceanic advection of temperature anomalies from the tropical Atlantic. Our results do not support the idea that wind stress curl anomalies associated with the atmospheric response pattern could be involved in such a feedback.

The marginal consistency of all the feedbacks we have identified suggests that if they play a role in any coupled ocean–atmosphere mode of variability, such a mode would not be robust against disruption by internal atmospheric variability. If there is an exception to this statement, it is most likely to be associated with the modulation of the trade winds, which we found to be a consistent feature of the atmospheric response in spring.

While we cannot rule out a coupled ocean–atmosphere mode of decadal variability, our results suggest

that the principal features of decadal variability in the North Atlantic Ocean will be explicable in terms of a passive ocean response to internal, quasi-stochastic, atmospheric variability (i.e., the ''uncoupled'' scenario we described in the introduction). We note again that this response involves ocean dynamics (e.g., Frankignoul et al. 1997) as well as ocean mixed-layer processes (Frankignoul and Hasselmann 1977) and that the possibility that a prefered timescale could develop in the ocean remains (e.g., Saravanan and McWilliams, 1998).

The absence of a coupled ocean–atmosphere mode of variability need not imply that there is no useful atmospheric predictability in the North Atlantic region arising from the influence of North Atlantic SST anomalies. On the contrary, we have identified a predictable atmospheric response that, at least in some regions and especially in spring, accounts for a significant fraction of the variance.[2] Put together with the evidence presented by Sutton and Allen (1997) that there is predictability in decadal fluctuations of North Atlantic SST, our results suggest that there may well be predictable aspects of decadal variability in the atmosphere over the North Atlantic. These aspects may include fluctuations in both the subtropical trade winds and in the higher latitude westerlies.

A key issue that requires further clarification is exactly which features in the North Atlantic SST distribution the atmosphere responds to. These, of course, are the features we would like to forecast. Is the atmospheric response induced by the full SST tripole pattern or merely a part of it? In particular, does the atmosphere respond especially to SST fluctuations in the Gulf Stream region? To address these issues requires further controlled experimentation with models, which we have begun to perform. Because of the possibility that experimentation with fixed SSTs may give a misleading response, experimentation with coupled models will be required as well.

Acknowledgments. We would like to thank Steve Jewson, Axel Timmermann, Anselm Grötzner, Mojib Latif, David Anderson, and Chris Folland for many helpful comments. We would also like to thank Simon Brown for his great help in making available to us the model data. Stephan Venzke was funded by the European Community via the SINTEX project and gratefully received travel support from the DAAD. MRA was supported by the NERC/RAL Ocean Dynamics SLA and an Advanced Research Fellowship from the UK-NERC. RTS was supported by the U.K. NERC through its Coupled Ocean Atmosphere Modelling Project and the U.K. Universities Global Atmospheric Modelling Programme. DPR was supported by the U.K. Public Meteorological Service Contract MSG-2/97. The AGCM integrations were performed at the Hadley Centre for Climate Prediction and Research by Alison Renshaw, John Davies, and David Sexton; data are subject to British Crown Copyright and have been reproduced with permission of the U.K. Met. Office. Computer time was provided by the U.K. Department of the Environment, Transport and Regions under Contract PECD 7/12/37.

[2] Note that these regions are mostly over the uninhabited ocean. However, as this paper focuses on ocean–atmosphere interactions, the response over the ocean is of particular interest.

REFERENCES

Allen, M. R., and L. A. Smith, 1996: Monte Carlo SSA: Detecting irregular oscillations in the presence of colored noise. *J. Climate,* **9,** 3373–3404.

——, and ——, 1997: Optimal filtering in singular spectrum analysis. *Phys. Lett.,* **234,** 419–428.

——, P. A. Stott, S. F. B. Tett, and W. J. Ingram, 1997: Noise in model predictions and optimal fingerprinting. *Eos, Trans. Amer. Geophys. Union,* **78** (Suppl.), 7126.

Barsugli, J. J., and D. S. Battisti, 1998: The basic effects of atmosphere–ocean thermal coupling on midlatitude variability. *J. Atmos. Sci.,* **55,** 477–493.

Basnett, T. A., and D. E. Parker, 1997: Development of the global mean sea level pressure data set GMSLP2. Hadley Centre, Climate Research Tech. Note 79, 52 pp.

Cullen, M. J. P., 1993: The unified forecast/climate model. *Meteor. Mag.,* **122,** 81–94.

da Silva, A. M., C. C. Young, and S. Levitus, 1994: *Atlas of Surface Marine Data.* NOAA/NESDIS, 83 pp.

Davies, J. R., D. P. Rowell, and C. K. Folland, 1997: North Atlantic and European seasonal predictability using an ensemble of multi-decadal AGCM simulations. *Int. J. Climatol.,* **17,** 1263–1284.

Deser, C., and M. L. Blackmon, 1993: Surface climate variations over the North Atlantic ocean during winter: 1900–1989. *J. Climate,* **6,** 1743–1753.

Frankignoul, C., and K. Hasselmann, 1977: Stochastic climate models. Part II: Application to sea-surface temperature variability and thermocline variability. *Tellus,* **29,** 284–305.

——, ——, and E. Zorita, 1997: A simple model of the decadal response of the ocean to stochastic wind forcing. *J. Phys. Oceanogr.,* **27,** 1523–1546.

Gates, W. L., 1992: AMIP: The atmospheric model intercomparison project. *Bull. Amer. Meteor. Soc.,* **73,** 1962–1970.

Grötzner, A., M. Latif, and T. P. Barnett, 1998: A decadal climate cycle in the North Atlantic Ocean as simulated by the ECHO coupled GCM. *J. Climate,* **11,** 831–847.

Hansen, D. V., and H. F. Bezdek, 1996: On the nature of decadal anomalies in North Atlantic sea surface temperature. *J. Geophys. Res.,* **101,** 8749–8758.

Hasselmann, K., 1976: Stochastic climate models. Part I: Theory. *Tellus,* **28,** 473–485.

——, 1979: On the signal-to-noise problem in atmospheric response studies. *Meteorology over the Tropical Oceans,* D. B. Shaw, Ed., Royal Meteorological Society, 251–259.

——, 1997: Multi-pattern fingerprint method for detection and attribution of climate change. *Climate Dyn.,* **13,** 601–611.

Hegerl, G. C., H. von Storch, K. Hasselmann, U. Cubasch, B. D. Santer, and P. D. Jones, 1996: Detecting greenhouse gas induced climate change with an optimal fingerprint method. *J. Climate,* **9,** 2281–2306.

Hurrell, J. W., 1995: Decadal trends in the North Atlantic Oscillation and relationships to regional temperature and precipitation. *Science,* **269,** 676–679.

James, I. N., and P. M. James, 1989: Ultralow-frequency variability in a simple circulation model. *Nature,* **342,** 53–55.

Kushnir, Y., and I. M. Held, 1996: Equilibrium atmospheric response to North Atlantic SST anomalies. *J. Climate,* **9,** 1208–1220.

Latif, M., and T. P. Barnett, 1994: Causes of decadal climate vari-

ability over the North Pacific and North America. *Science*, **266,** 634–637.

Lean, J., J. Beer, and R. Bradley, 1995: Reconstruction of solar irradiance since 1600: Implications for climate change. *Geophys. Res. Lett.*, **22,** 3195–3198.

Lorenz, E. N., 1956: Empirical orthogonal functions and statistical weather prediction. Scientific Rep. 1, Department of Meteorology, Massachusetts Institute of Technology, 49 pp. [Available from MIT, 77 Massachusetts Ave., Cambridge, MA 02139.]

——, 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20,** 130–141.

——, 1990: Can chaos and intransitivity lead to interannual variability. *Tellus*, **42A,** 378–389.

Palmer, T. N., and Z. Sun, 1985: A modelling and observational study of the relationship between sea surface temperature in the northwest Atlantic and the atmospheric general circulation. *Quart. J. Roy. Meteor. Soc.*, **111,** 947–975.

——, and D. L. T. Anderson, 1994: The prospects for seasonal forecasting—A review paper. *Quart. J. Roy. Meteor. Soc.*, **120,** 755–793.

Parker, D. E., C. K. Folland, A. Bevan, M. N. Ward, M. Jackson, and F. Maskell, 1995: Marine surface data for analysis of climate fluctuations on interannual to century timescales. *Natural Climate Variability on Decadal to Century Time scales*, D. G. Martinson et al., Eds., National Academy Press., 241–250.

Peng, S., W. A. Robinson, and M. P. Hoerling, 1997: The modeled atmospheric response to midlatitude SST anomalies and its dependence on background circulation states. *J. Climate*, **10,** 971–987.

Phillips, T. J., 1994: A summary documentation of the AMIP models. PCMDI Report No. 18, Lawrence Livermore National Laboratory, Livermore, CA. [Avialable from Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, CA 94550-9234.]

Pitcher, E. J., M. L. Blackmon, and S. Muñoz, 1988: The effect of North Pacific temperature anomalies on the January climate in a general circulation model. *J. Atmos. Sci.*, **45,** 173–188.

Renshaw, A. C., D. P. Rowell, and C. K. Folland, 1998: ENSO responses and low frequency weather variability in the North Pacific/American sector 1943–93. *J. Climate*, **11,** 1073–1093.

Robock, A. D., and J. Mao, 1995: A volcanic signal in surface temperature observations. *J. Climate*, **8,** 1086–1103.

Rogers, J. C., 1990: Patterns of low frequency monthly sea level pressure variability (1899–1986) and associated wave cyclone frequencies. *J. Climate*, **3,** 1364–1379.

Rowell, D. P., 1998: Assessing potential seasonal predictability with an ensemble of multidecadal GCM simulations. *J. Climate*, **11,** 109–120.

——, C. K. Folland, K. Maskell, and M. N. Ward, 1995: Variability of summer rainfall over tropical north Africa (1906–92): Observations and modelling. *Quart. J. Roy. Meteor. Soc.*, **121,** 669–704.

Saravanan, R., and J. C. McWilliams, 1998: Advective ocean–atmosphere interaction: An analytical stochastic model with implications for decadal variability. *J. Climate*, **11,** 165–188.

Sutton, R. T., and M. R. Allen, 1997: Decadal predictability of North Atlantic sea surface temperature and climate. *Nature*, **388,** 563–567.

Thacker, W. C., and R. Lewandowicz, 1996: Climate indices, principal components, and the Gauss–Markov theorem. *J. Climate*, **9,** 1942–1958.

White, W. B., J. Lean, D. R. Cayan, and M. D. Dettinger, 1997: Response of global upper ocean temperature to changing solar irradiance. *J. Geophys. Res.*, **102,** 3255–3266.

Willmott, C. J., C. M. Rowe, and Y. Minz, 1985: A global archive of land cover and soils data for use in general circulation models. *J. Climatol.*, **5,** 119–143.

Zhang, Q., M. Brasseville, and A. Beneviste, 1994: Early warning of slight changes in systems. *Automatica*, **30,** 95–113.