Using Neutral Singular Vectors to Study Low-Frequency Atmospheric Variability

JASON C. GOODMAN AND JOHN MARSHALL

Program in Atmospheres, Oceans and Climate, Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts

(Manuscript received 6 September 2001, in final form 10 May 2002)

ABSTRACT

The authors explore the use of the ''neutral vectors'' of a linearized version of a global quasigeostrophic atmospheric model with realistic mean flow in the study of the nonlinear model's low-frequency variability. Neutral vectors are the (right) singular vectors of the linearized model's tendency matrix that have the smallest eigenvalues; they are also the patterns that exhibit the largest response to forcing perturbations in the linear model. A striking similarity is found between neutral vectors and the dominant patterns of variability (EOFs) observed in both the full nonlinear model and in the real world. The authors discuss the physical and mathematical connection between neutral vectors and EOFs.

Investigation of the ''optimal forcing patterns''—the left singular vectors—proves to be less fruitful. The neutral modes have equivalent barotropic vertical structure, but their optimal forcing patterns are baroclinic and seem to be associated with low-level heating. But the horizontal patterns of the forcing patterns are not robust and are sensitive to the form of the inner product used in the singular vector decomposition analysis. Additionally, applying ''optimal'' forcing patterns as perturbations to the full nonlinear model does not generate the response suggested by the linear model.

1. Introduction

a. Background

A large body of observational evidence documents the presence of preferred modes of variability of the midlatitude atmosphere (see, e.g., Walker and Bliss 1932; Barnston and Livezey 1987; Cayan 1992; Kushnir 1994; Hurrell 1995; White and Peterson 1996; Thompson and Wallace 1998; Czaja and Marshall 2001). A relatively small number of patterns explain a large fraction of the variability of atmospheric flow on timescales from weeks to decades. The most well known of these are the North Atlantic Oscillation (NAO), a dipole mode in the North Atlantic (Walker and Bliss 1932), and the Pacific–North America pattern (PNA), extending from the tropical Pacific to North America (Wallace et al. 1993). Patterns that are receiving increasing attention are the zonally symmetric ''annular modes'' of Thompson and Wallace (1998) in the Northern and Southern Hemispheres.

In this study we investigate the dynamical nature of these modes of variability by developing a description that is rooted in dynamics rather than statistics. We continue a line of enquiry begun by Marshall and Molteni (1993) and Navarra (1993), who associate preferred modes of variability with ''neutral vectors''—singular vectors of the tendency matrix of a linearized model that have the smallest singular values.

Our work extends that of Marshall and Molteni by considering both the neutral vectors and their forcing, and by studying their connection to empirical orthogonal functions (EOFs). We expand on Navarra's work by using a model that includes baroclinic dynamics and that is capable of producing patterns similar to those seen in nature.

b. Introduction to neutral vectors

Marshall and Molteni (hereafter MM) were interested in atmospheric patterns that tended to persist in a given state for long periods of time. They attempted to compute patterns of maximum persistence by beginning with a forced three-layer quasigeostrophic potential vorticity (QGPV) model, which we write as

$$\frac{\partial}{\partial t}\boldsymbol{\Psi} = \mathcal{M}(\boldsymbol{\Psi}) + \mathbf{f},$$

where $\boldsymbol{\Psi}$ is a vector representing the model streamfunction, $\mathcal{M}$ is a nonlinear tendency operator, and $\mathbf{f}$ is a potential vorticity (PV) source term. Linearizing this model about some basic state gives

$$\frac{\partial}{\partial t}\boldsymbol{\Psi} = \mathbf{M}\boldsymbol{\Psi} + \mathbf{f}, \tag{1}$$

Corresponding author address: Jason C. Goodman, Department of Geophysical Sciences, University of Chicago, 5734 S. Ellis Ave., Chicago, IL 60640.
E-mail: goodmanj@uchicago.edu

where $\mathbf{M}$ is a tendency matrix, $\boldsymbol{\Psi}$ now represents streamfunction perturbations upon some specified basic state, and $\mathbf{f}$ is now a perturbation in the steady-state model forcing field.

MM were interested in the free, unforced perturbations that displayed the smallest time tendency. Free, unforced waves obey

$$\frac{\partial}{\partial t}\boldsymbol{\Psi} = \mathbf{M}\boldsymbol{\Psi}. \tag{2}$$

To find the modes with the smallest time tendency, MM minimized the expression

$$\lambda^2 = \frac{\left\langle \frac{\partial}{\partial t}\boldsymbol{\Psi}, \frac{\partial}{\partial t}\boldsymbol{\Psi} \right\rangle}{\langle \boldsymbol{\Psi}, \boldsymbol{\Psi} \rangle}, \tag{3}$$

where $\langle \mathbf{a}, \mathbf{b} \rangle$ is an inner product of $\mathbf{a}$ and $\mathbf{b}$. This expression minimizes the size of the mode's tendency, normalized by the magnitude of the mode itself. Inserting (2) into (3), the minimization condition can be rewritten

$$\lambda^2 = \frac{\langle \mathbf{M}\boldsymbol{\Psi}, \mathbf{M}\boldsymbol{\Psi} \rangle}{\langle \boldsymbol{\Psi}, \boldsymbol{\Psi} \rangle} = \frac{\langle \mathbf{M}^\dagger\mathbf{M}\boldsymbol{\Psi}, \boldsymbol{\Psi} \rangle}{\langle \boldsymbol{\Psi}, \boldsymbol{\Psi} \rangle},$$

where $\mathbf{M}^\dagger$ is the adjoint of $\mathbf{M}$. The $\boldsymbol{\Psi}$ that minimize λ will be the eigenvectors $\boldsymbol{\Psi}_n$ of $\mathbf{M}^\dagger\mathbf{M}$ with minimum eigenvalue λ_n^2:

$$\lambda^2 = \frac{\langle \mathbf{M}^\dagger\mathbf{M}\boldsymbol{\Psi}_n, \boldsymbol{\Psi}_n \rangle}{\langle \boldsymbol{\Psi}_n, \boldsymbol{\Psi}_n \rangle} = \frac{\langle \lambda_n^2\boldsymbol{\Psi}_n, \boldsymbol{\Psi}_n \rangle}{\langle \boldsymbol{\Psi}_n, \boldsymbol{\Psi}_n \rangle}.$$

These $\boldsymbol{\Psi}_n$ are the neutral vectors.

What are the dynamics of a neutral vector? The matrix $\mathbf{M}$ encapsulates the Rossby wave propagation, downstream advection, and dissipation terms of the QGPV equation. For $\mathbf{M}\boldsymbol{\Psi}$ to be small, dissipation must be weak (implying large-scale patterns), and there must be a near balance between the propagative and advective terms.

Suppose we took the model in (1) and looked at the forced, linear, stationary response to a thermal forcing $\mathbf{f}$:

$$0 = \mathbf{M}\boldsymbol{\Psi} + \mathbf{f}. \tag{4}$$

What pair of forcing and response will have the largest response per unit forcing? We want to find the $\boldsymbol{\Psi}$ and $\mathbf{f}$ that will maximize

$$\lambda^{-2} = \langle \boldsymbol{\Psi}, \boldsymbol{\Psi} \rangle / \langle \mathbf{f}, \mathbf{f} \rangle.$$

Since $\mathbf{f} = -\mathbf{M}\boldsymbol{\Psi}$, this is equivalent to *minimizing*

$$\lambda^2 = \frac{\langle \mathbf{M}\boldsymbol{\Psi}, \mathbf{M}\boldsymbol{\Psi} \rangle}{\langle \boldsymbol{\Psi}, \boldsymbol{\Psi} \rangle}.$$

This is exactly the condition satisfied by the neutral vectors. Thus, the neutral vectors are not only the most stationary modes in the unforced time-evolving model, they are also the forced, stationary modes that exhibit the largest response to external forcing. Such a strong response to forcing is necessary if a mutually coupled atmosphere–ocean interaction is to occur.

Interestingly, this means we cannot only find the neutral vectors, $\boldsymbol{\Psi}_n$, but also the ''optimal forcing patterns'' $\mathbf{f}_n$ that maximally excite them, by solving $\mathbf{M}\boldsymbol{\Psi}_n + \mathbf{f}_n = 0$. The $\boldsymbol{\Psi}$ are the right singular vectors of $\mathbf{M}$; the $\mathbf{f}_n$ are the left singular vectors.

The linearized system (1) posesses unstable growing modes. In addition, the full model dynamics is strongly nonlinear. One may wonder whether these complications render the neutral vector analysis moot. We address this issue in section 5. By studying the dynamics of simple unstable nonlinear systems, we demonstrate that the linear analysis described above remains relevant in these systems.

Singular vectors are often applied to investigate the role of nonnormal growth in the production and maintenance of transient perturbations to the atmosphere (beginning with the study of Farrell 1982), and their interactions with the mean flow (Whitaker and Sardeshmukh 1998). A related problem considers the problem of optimal excitation patterns; see Farrell (1989) and Molteni and Palmer (1993). The linear version of the Principal Interaction Pattern (PIP)/Principal Oscillation Pattern (POP) techniques introduced by Hasselmann (1988) also use singular vectors to find patterns that best describe a system's tendency. The technique has also been used to find optimal amplification patterns in classical, nonrotating fluid dynamics (Andersson et al. 1999; Luchini 2000). However, in all the aforementioned applications the focus is on the singular vectors with the *largest* singular values, the most rapidly evolving and changing modes. Here our emphasis is on the physical relevance of the singular vectors with the *smallest* singular values.

c. Goals of this study

In this work we study neutral vectors and their connection EOFs using the three-layer quasigeostrophic (QG) model of MM, which was designed to produce mean flows and patterns of variability resembling those found in nature. The principal questions to be answered are the following.

1) Is there a connection between the leading modes of variability of the model and the neutral modes of a linearized version of the model?
2) Why should such a connection exist?
3) Do the model's neutral modes resemble observed modes of strong variability?
4) What forcing patterns optimally excite these modes?

An affirmative answer to the first question establishes neutral vectors as a useful tool for studying model variability. MM gave a rough answer to this question, inspiring the more detailed analysis presented here. An affirmative answer to the third question establishes neutral vectors as a useful tool for studying observed patterns of variability. The last question is particularly rel-

evant to understanding coupled modes. If our model can identify, for example, an SST forcing pattern that strongly excites a pattern that resembles an observed mode of variability, then it becomes more likely that SST can drive interannual variability of this mode, and we can check for correlations between amplitude of the SST forcing pattern and of the response. The role of neutral vectors in air–sea interaction is studied in detail in Goodman and Marshall (2002, hereafter GM02).

The three-layer QG model used here lacks some of the dynamics and physics of a full primitive equation GCM, but provides a useful and computationally inexpensive tool to develop and test our procedure. The analysis developed here can easily be applied to a full GCM, with additional (but not prohibitive) computational cost. We chose the three-layer model mainly because the model and its linearized version are well developed; in addition, a three-layer model enables us to investigate the vertical structure of heating that excites a model response. In a two-layer model, temperature is defined at the single layer interface, and one cannot distinguish between low- and high-level heating.

The model, neutral vector, and EOF computations are described in section 2. The neutral vectors are discussed in section 3, and the connection between neutral vectors and EOFs is discussed in section 3b. We consider the model's optimal forcing patterns in section 4, including an attempt to discover whether the optimal forcing patterns remain optimal in the full nonlinear model in section 4c. Discussion of the relevance of our linear approach to nonlinear dynamical systems is found in section 5.

2. Molteni's three-layer quasigeostrophic model

Molteni's model is discussed in some detail in his thesis (Molteni 1994) and in MM. The model is a three-layer, global, spectral model at T21 resolution, with pressure as a vertical coordinate. The QGPV equations are discretized at the three pressure levels (200, 500, and 800 mb), giving prognostic equations for PV:

$$\frac{\partial}{\partial t}q_1 = -J(\psi_1, q_1) - D_1(\psi_1, \psi_2) + S_1,$$

$$\frac{\partial}{\partial t}q_2 = -J(\psi_2, q_2) - D_2(\psi_1, \psi_2, \psi_3) + S_2,$$

$$\frac{\partial}{\partial t}q_3 = -J(\psi_3, q_3) - D_3(\psi_1, \psi_2) + S_3.$$

The ψ_n and q_n are the streamfunction and QGPV at each level, D_n encapsulate various linear dissipative processes (see Molteni 1994), and the S_n are a constant PV source term. PV is defined as:

$$q_1 = \nabla^2\psi_1 - R_1^{-2}(\psi_1 - \psi_2) + f_{\text{cori}}$$

$$q_2 = \nabla^2\psi_2 + R_1^{-2}(\psi_1 - \psi_2) - R_2^{-2}(\psi_2 - \psi_3) + f_{\text{cori}}$$

$$q_3 = \nabla^2\psi_3 + R_2^{-2}(\psi_2 - \psi_3) + f_{\text{cori}}(1 + h/H_0),$$

where R_1 and R_2 are (spatially constant) Rossby radii of deformation, h is the height of topography, H_0 is a topographic scaling factor, and f_{cori} is the Coriolis parameter.

The model covers the entire globe, so it does not have artificial ''walls'' at the equator that may spuriously reflect planetary waves. However, since QG dynamics is not really appropriate near the equator (and the constant values chosen for the Rossby radii are certainly inappropriate there), the model's behavior in the Tropics should not be taken too seriously. In the midlatitudes, however, the model's dynamics are more reasonable, and it can produce a very good mean flow field through careful specification of the constant forcing fields S_n. This is done by setting the S_n equal to the opposite of the average PV tendencies obtained by inserting observed streamfunction fields into a version of the model equations from which the S_n are omitted. This technique has been found to give the model a stable climatology that is near the observed fields used to generate the S_n. The technique is similar to the ''flux correction'' used to eliminate climate drift in coupled GCMs. The S_n used for this study are computed from European Centre for Medium-Range Weather Forecasts streamfunction analyses, using data from December through March for 1983 through 1993 (Michelangeli and Vautard 1998). As a result, the model attempts to simulate the Northern Hemisphere winter climatology.

Figure 1 presents mean fields computed from a 5000-day integration of the model. Comparisons with observed wintertime mean streamfunction (not shown) show that the model does a very good job at reproducing the mean flow of the observations used to compute *S.* The midlatitude jets are of realistic strength, and their pattern of confluence/diffluence relative to the underlying continents and oceans is captured. Generally, the model's mean state differs from observations by only a few percent. Eddy activity, as shown by streamfunction standard deviation in the lower panel, is less well captured in the model. Eddy activity is of the correct magnitude, with storm tracks over the Northern Hemisphere oceans, much as observed. The Pacific storm track is captured rather well. However, the modeled Atlantic storm track does not have a sharp northern boundary, and eddy activity over the pole is somewhat stronger than observed.

We use this 5000-day run to compute empirical orthogonal functions (North 1984) of the model output. The EOFs are computed over the entire model domain, using monthly mean data. EOFs computed with weekly averaged data were found to be quite similar.

a. Computing neutral vectors

The linearized form of the model equations is

$$\frac{\partial}{\partial t}q = -J(\overline{\psi}, q) - J(\psi, \overline{q}) - D(\psi) + S, \qquad (5)$$

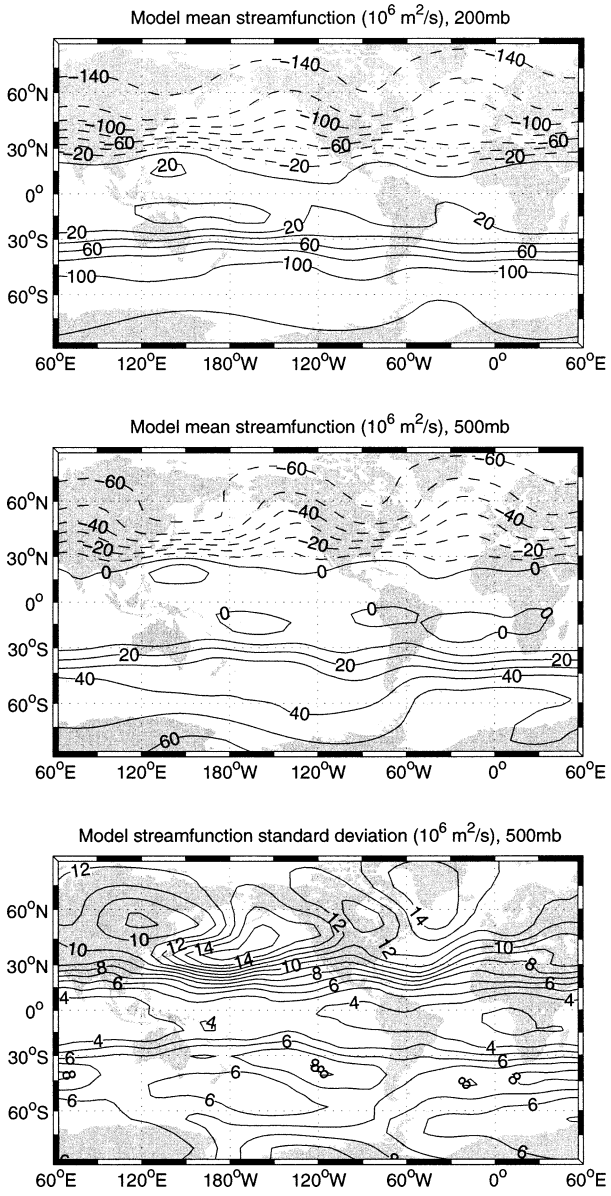

FIG. 1. Three-layer QG model mean climatology, based on 5000 daily fields. Top panel: mean streamfunction, 200 mb, contour interval 20×10^6 m^2 s^{-1}. Middle panel: mean streamfunction, 500 mb, contour interval 10×10^6 m^2 s^{-1}. Lower panel: Model streamfunction standard deviation, 500 mb, contour interval 10^6 m^2 s^{-1}.

where overlined terms represent the basic state and unmarked terms represent perturbation quantities. Note the presence of a linearized PV forcing anomaly S. The nature of this forcing is left unspecified: it could be a sea surface temperature anomaly, anthropogenic forcing, or some other source. We perform a PV inversion to express the equation in units of streamfunction tendency:

$$\frac{\partial}{\partial t}\psi = -Q^{-1}[J(\overline{\psi}, q)) + J(\psi, \overline{q}) + D(\psi)] + f, \quad (6)$$

where $f = Q^{-1}(S)$ is the streamfunction forcing field. We now discretize and write this in vector form thus:

$$\frac{\partial}{\partial t}\mathbf{\Psi} = \mathbf{M}\mathbf{\Psi} + \mathbf{f}. \quad (7)$$

Here, $\mathbf{M}$ is the streamfunction tendency operator, and $\mathbf{f}$ is the streamfunction forcing perturbation. We now minimize the tendency (for the unforced problem) or the forcing (for the steady problem)

$$\lambda^2 = \frac{\left\langle \frac{\partial}{\partial t}\mathbf{\Psi}\frac{\partial}{\partial t}\mathbf{\Psi} \right\rangle}{\langle \mathbf{\Psi}, \mathbf{\Psi} \rangle} \quad \text{unforced, time evolving;} \quad (8)$$

$$\lambda^2 = \frac{\langle \mathbf{f}, \mathbf{f} \rangle}{\langle \mathbf{\Psi}, \mathbf{\Psi} \rangle} \quad \text{forced, steady;} \quad (9)$$

by computing the left and right singular vectors of $\mathbf{M}$ as described in section 1b. The λ may be interpreted as inverse timescales; the corresponding mode will remain relatively unchanged for a time $\leq \lambda^{-1}$.

One can make various choices for the inner product in Eq. (3); this boils down to the question, ''We want to find the patterns whose tendency is smallest. . . but what do we mean by smallest?'' To begin, we choose an inner product identical to that used by MM: the norm of a streamfunction vector is chosen to be proportional (in our spherical domain) to the kinetic energy of the flow. We refer to this as the ''KE norm''; a different choice will be considered in section 4a.

The only remaining difficulty is to construct the $\mathbf{M}$ matrix. This job is made much simpler by the existence of a linearized version of the Molteni model code. We linearize about the mean state of the 5000-day run shown in Fig. 1. Using this linearized code, we simply compute the tendencies of a complete set of orthonormal spectral Green's function perturbations and use those tendencies to build up an explicit matrix $\mathbf{M}$, column by column. That is, we build a matrix whose first column contains the normalized response of the linearized model to a perturbation of the model's first spherical harmonic, whose second column contains the response to perturbation along the second harmonic, and so on for all 1518 degrees of freedom. The eigenvalues and eigenvectors of this matrix are then computed using Matlab's built-in dense matrix eigensolver.

b. Empirical orthogonal functions

The first few EOFs of the model's monthly mean streamfunction fields are shown in Fig. 2. The EOFs are different from those described by, for example, Molteni et al. (1988); this is probably because Molteni et al. computed hemispheric EOFs of the *eddy* fields (i.e., zonal mean components are removed), while here global EOFs of the full streamfunction are computed. Also, Molteni's patterns are EOFs of observed fields, while we compute EOFs of the model output. Nonetheless,

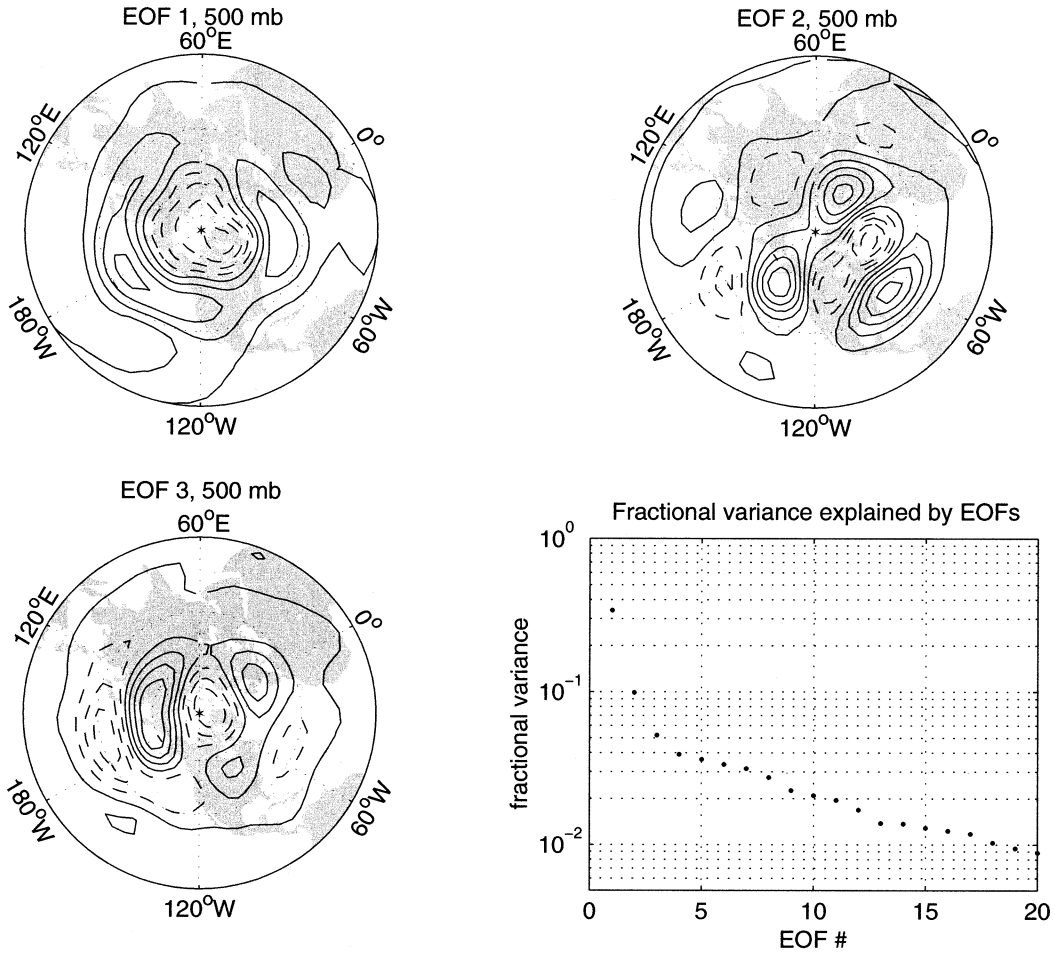

FIG. 2. Empirical orthogonal functions computed from monthly means of a 5000-day integration of the Molteni QG model. EOFs are computed on global model output at all three levels; EOF amplitude at 500 mb is shown here.

the model EOFs do resemble observations. The first EOF, explaining 35% of the variance of monthly means, has the dipolar nature of the NAO in the Atlantic, but is much more zonally extensive. It more closely resembles an annular mode, or the ''Arctic Oscillation'' (AO; Thompson and Wallace 1998), which are commonly seen in full-hemisphere EOF analyses of observations and models. Wallace (2000) argues that the NAO and the AO are really the same phenomenon: an annular mode that is somewhat stronger in the Atlantic than elsewhere.[1]

The second EOF of the model strongly resembles the planetary wave pattern associated with the PNA pattern (Wallace et al. 1993). We see a wave train extending from the subtropical Pacific near the date line over the Gulf of Alaska and Canada, and ending in the subtropical west Atlantic.

The third EOF displays a wave train extending from the midlatitude Pacific west of the dateline over the pole to western Europe and has wavelengths similar to the first two EOFs.

3. Neutral vectors

When we computed neutral vectors for the three-layer model in the manner discussed in section 2a, patterns were found in both hemispheres, whereas the EOFs were confined to the Northern Hemisphere. The reason for this difference will be discussed in section 3b. We feel it is unlikely that a dynamical connection exists between wave patterns in alternate hemispheres; it is more likely that two separate neutral patterns in the two hemispheres share similar eigenvalues. The singular vector decomposition (SVD) analysis cannot distinguish between two modes with similar eigenvalues and will return two arbitrary orthogonal linear superpositions of the two

[1] We note that papers that use EOFs localized to the Atlantic domain (Cayan 1992, e.g.) tend to find a more localized NAO pattern, while hemispheric EOF calculations (Thompson and Wallace 1998) tend to find the annular AO pattern. A possible reason for this is given in appendix B of Goodman (2001).

modes. A ''rotated neutral vector analysis,'' along the lines of rotated EOFs (Richman 1986), might help to separate the modes.

Instead of rotating the modes, we focus on the Northern Hemisphere by adding an artificial damping term to the **M** matrix. This term is proportional to $\sin(\phi/2 - \pi/4)^6$ (where ϕ is latitude in radians) and damps PV anomalies with a timescale of 5 days at the South Pole, 8 days at 45°S, 40 days at the equator, and 1500 days at 45°N. Thus, any otherwise neutral mode in the Southern Hemisphere will have a significant tendency due to this damping effect, making it less neutral. The effect is to reorder the neutral mode patterns, giving preference to Northern Hemisphere modes.

a. Neutral vector structure

In Figs. 3–5, we show neutral vectors and optimal forcing patterns computed using the ''kinetic energy'' norm used by MM. We describe the neutral vectors (right-hand column) in detail in this section; the optimal forcing patterns are discussed in section 4a.

The first neutral vector (Fig. 3, right column) shows a roughly zonally symmetric pattern, with a negative center over the pole surrounded by a positive annulus near 50°N. This annulus has enhanced energy over western Siberia and the North Atlantic. The whole pattern is equivalent barotropic and broadly resembles the first model EOF, although its midlatitude annulus is farther north than in EOF 1.

The second neutral vector shows a similar banded structure, but the maxima and minima are shifted southward relative to neutral vector 1. The positive annulus for this pattern lies at 40°N, with a minimum at 70°. Once again, the mode is barotropic, with amplitude increasing with height. The projection of this pattern onto the NAO pattern is quite strong.

The third neutral vector displays a pattern resembling the PNA pattern. We see a barotropic wave train extending from the subtropical Pacific over the Gulf of Alaska, northern Canada, and into the subtropical east Atlantic. The subtropical Pacific maximum is shifted west across the date line compared to the observed PNA, but elsewhere, the resemblance between neutral vector 3, EOF 2, and the PNA is very strong.

It should be clear that these patterns strongly resemble the model EOFs and observed patterns of low-frequency variability. To emphasize this point, and to demonstrate that these resemblances are not the result of chance correlations in a low-dimensional system, Fig. 6 shows the cross correlations between the shapes of the model EOFs (Fig. 2) and the neutral vectors (Figs. 3–5). Correlation is computed over all vertical levels. The absolute value of the correlation for each pair is indicated by the size of the dots. Correlations less than 0.2 are not plotted.

The observations discussed above are borne out in this figure: neutral vectors 1 and 2 project onto EOF 1, and neutral vector 3 projects onto EOF 2. We also see that neutral vector 1 projects onto EOF 3. But more importantly, the strongest correlations are among the first few EOFs and neutral vectors (the large dots are clustered in the upper left corner). If the correlations arose by chance, we would expect this figure to show a random scattering of small points throughout the domain.

We have visually compared these patterns with the results of d'Andrea (2002), who sought patterns which minimize the tendency of a *nonlinear* model, using a method of steepest descent. We find good agreement between some of the neutral vectors shown here and some of d'Andrea's patterns. Interestingly, in some cases, the neutrality of a mode in the nonlinear model is sign dependent: the ''positive'' phase of a neutral vector may appear in d'Andrea's set of patterns with minimal tendency, while the ''negative'' phase may not.

Since the first few neutral vectors project strongly onto the EOFs, we should expect that a substantial amount of the model's natural variability resides in the subspace of the first few neutral vectors. Figure 7 demonstrates that this is indeed the case. The three most neutral vectors explain more of the variance of a 5000-day model run than any other singular vector. The first three EOFs explain 50% of the variance of monthly means, while the first three neutral vectors explain 37% of the variance—almost as much.[2]

Many papers have been written that project observed variability onto the first few EOFs to study its statistics (Hannachi 1997; Haines and Hannachi 1995; Molteni et al. 1988), or that use EOFs to generate a reduced-subspace model that encapsulates most of the system's variability (Achatz and Branstator 1999; Kaplan et al. 2000). We find here that neutral vectors are almost equally good for these purposes. They also have the advantage over EOFs that they represent dynamically important modes of the system, rather than being empirically selected.

b. Relationship between neutral vectors and EOFs

In the previous section, we observed a close connection between EOFs and neutral vector patterns. We now explore the reasons for this connection. Some of the results presented in this section have been published by Navarra (1993), but because our discussion of the importance of the form of the $\mathbf{f}\mathbf{f}^{\dagger}$ matrix goes beyond Navarra's study, we present the derivation in full here.

Up to this point, we have ignored the effects of the nonlinear synoptic eddies. Synoptic systems can create a time mean eddy potential vorticity flux that drives the mean flow. Variations in this flux can be incorporated in our linearized system through the external forcing term f.

[2] Since EOFs, by definition, maximize explained variance, the variance explained by neutral vectors must be smaller.

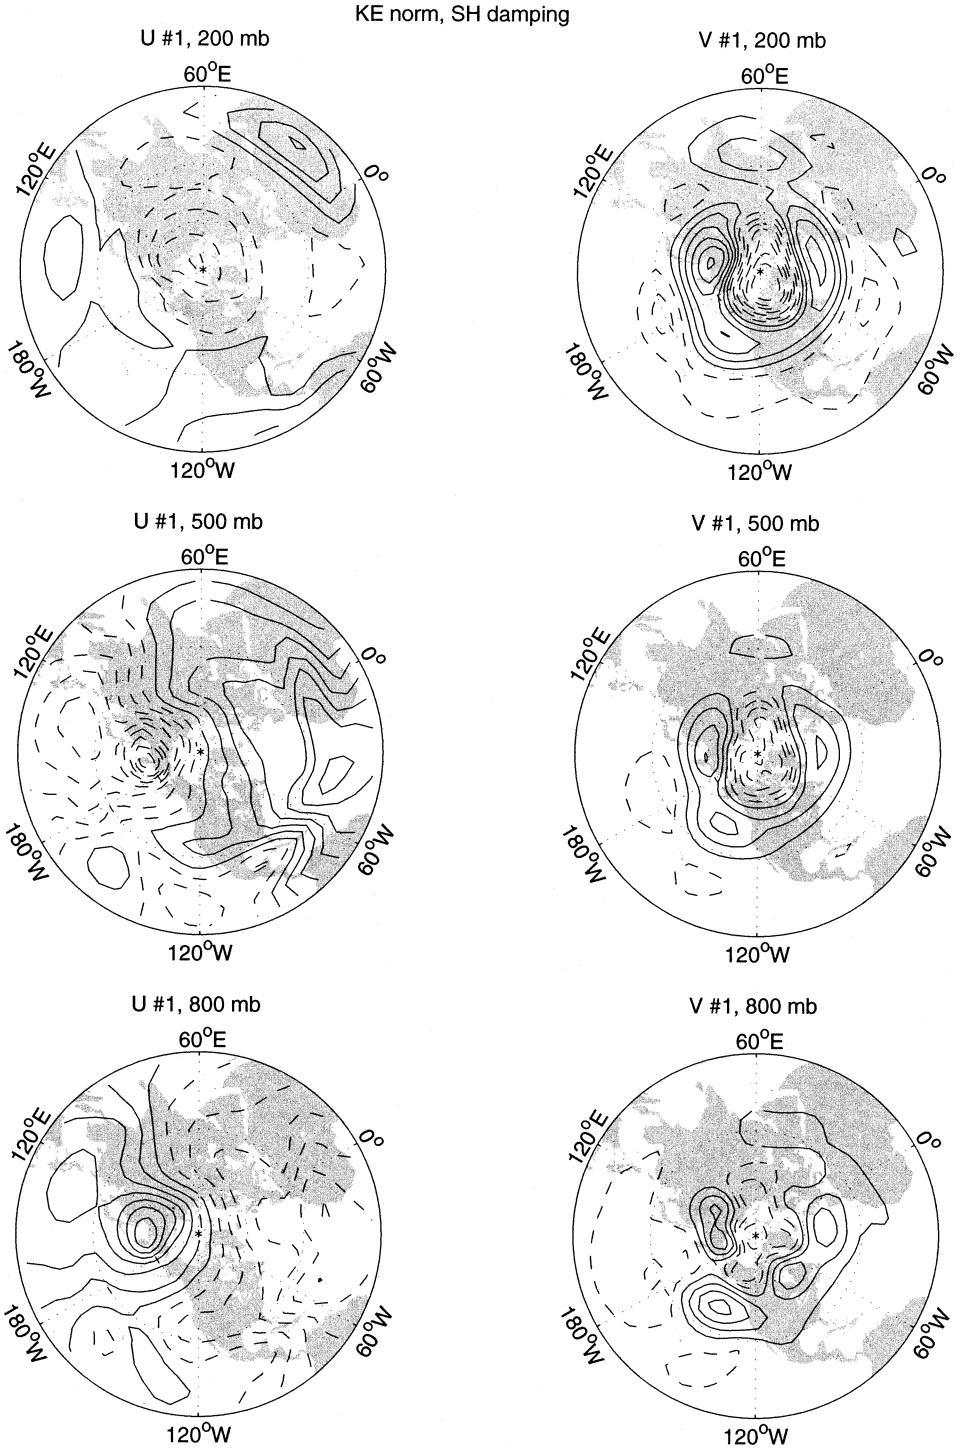

FIG. 3. Right column: Neutral vector 1 for model climatology, KE inner product. Left column: corresponding optimal forcing pattern. Contour interval is arbitrary, but consistent from level to level for *U*; contour spacing is halved for *V* at 800 mb to show details of near-surface structure.

Suppose the atmosphere consists of a number of ''slow modes'' (stationary planetary wave patterns like the NAO and PNA), which respond in a linear way to stochastic forcing generated by ''fast modes'' (transient eddy forcing), such that during any observation period n, the stochastic forcing $\mathbf{f}_n$ excites a planetary wave response $\mathbf{\Psi}_n$. We assume that the observation period is sufficiently long that initial transients die out, and a

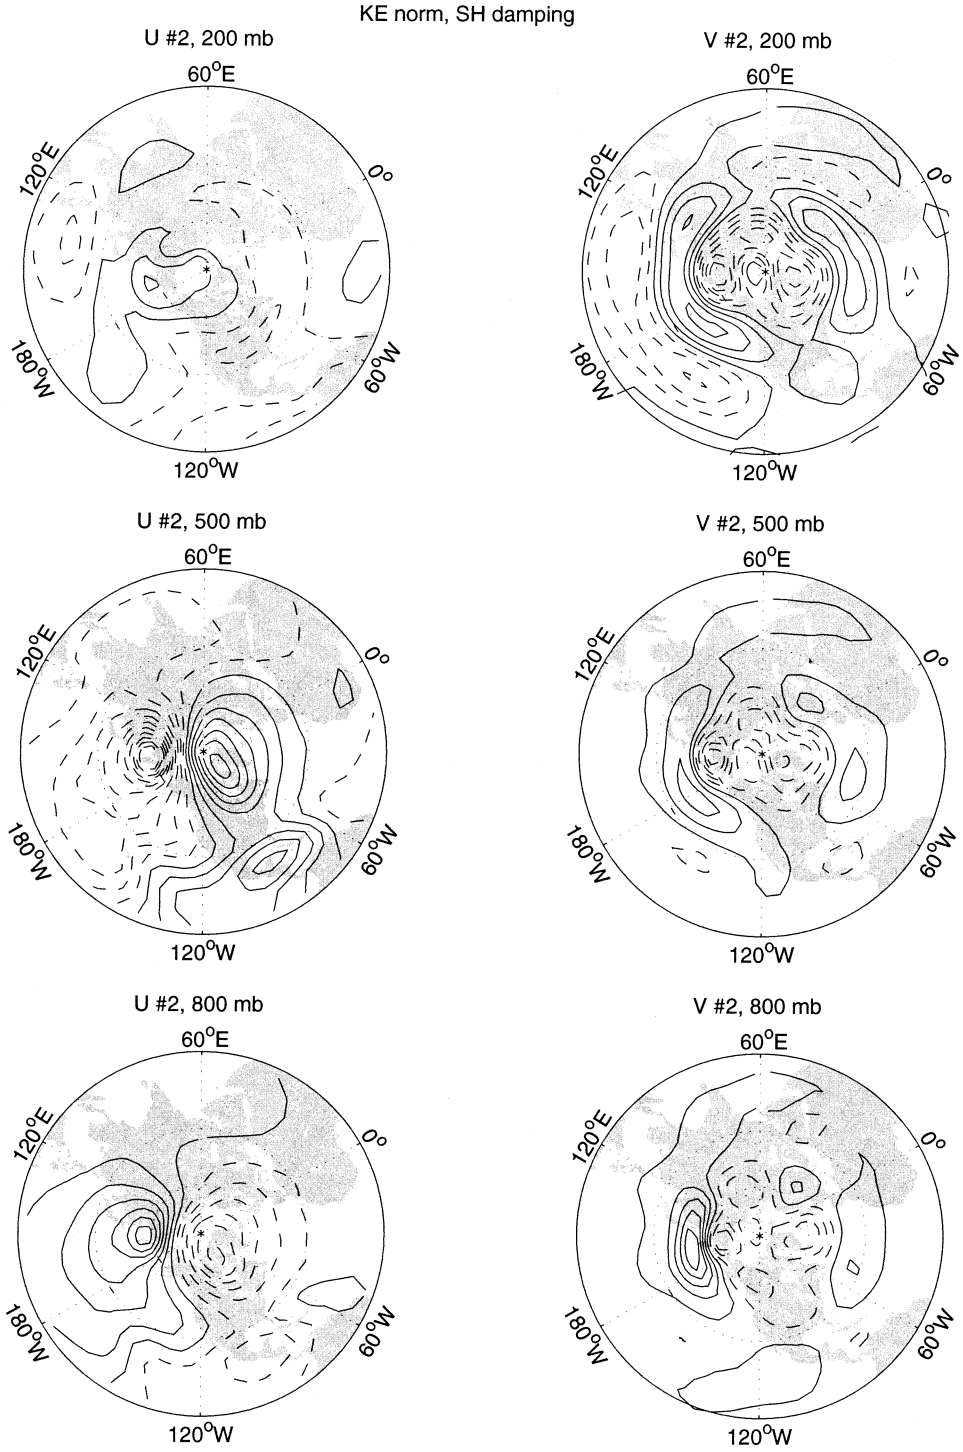

FIG. 4. Same as Fig. 3, but for neutral vector 2. Right column: neutral vector; left column: corresponding optimal forcing pattern.

steady-state balance exists between forcing and response:

$$\mathbf{M}\boldsymbol{\Psi}_n = \mathbf{f}_n \rightarrow \boldsymbol{\Psi}_n = \mathbf{M}^{-1}\mathbf{f}_n,$$

or, defining matrices $\boldsymbol{\Psi}$ and $\mathbf{f}$ whose columns are the $\boldsymbol{\Psi}_n$ and $\mathbf{f}_n$:

$$\boldsymbol{\Psi} = \mathbf{M}^{-1}\mathbf{f}.$$

The EOFs are defined as the eigenvectors of the cozvariance matrix $\boldsymbol{\Psi}\boldsymbol{\Psi}^{\dagger}$:

$$\boldsymbol{\Psi}\boldsymbol{\Psi}^{\dagger} = \mathbf{M}^{-1}\mathbf{f}\mathbf{f}^{\dagger}\mathbf{M}^{-\dagger},$$

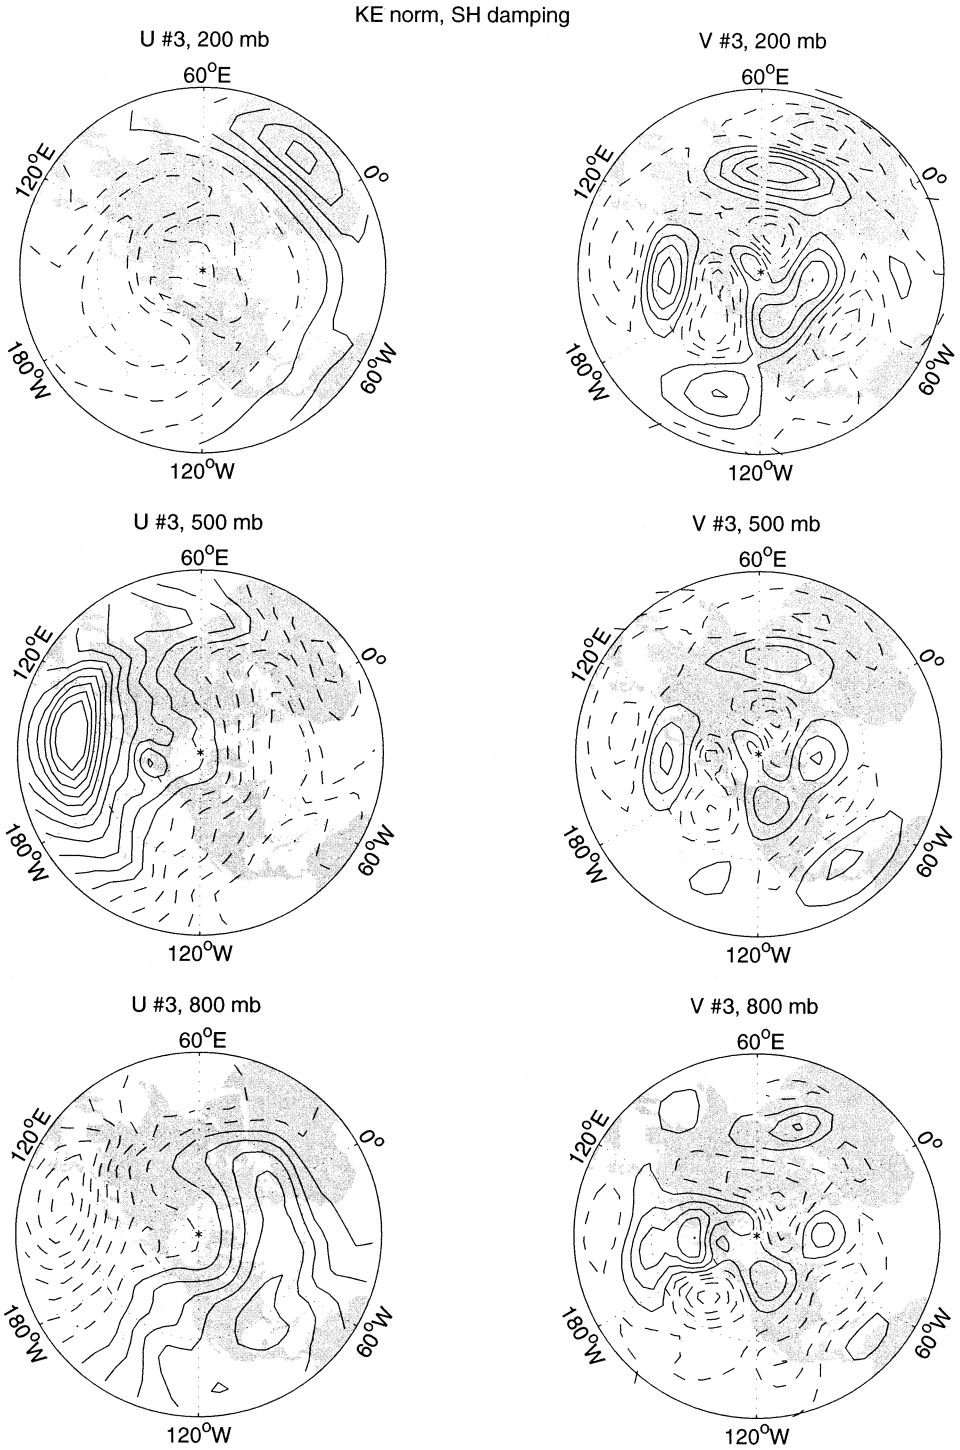

FIG. 5. Same as Fig. 3, but for neutral vector 3. Right column: neutral vector; left column: corresponding optimal forcing pattern.

where $\mathbf{M}^{-\dagger}$ is shorthand for $(\mathbf{M}^{\dagger})^{-1}$. Consider for a moment the case where the eddy forcing is isotropic (uniform variance everywhere) and spatially uncorrelated. If $\mathbf{f}$ has unit amplitude, then $\mathbf{f}\mathbf{f}^{\dagger} = \mathbf{I}$ so that

$$\mathbf{\Psi}\mathbf{\Psi}^{\dagger} = \mathbf{M}^{-1}\mathbf{M}^{-\dagger}. \tag{10}$$

We may decompose the $\mathbf{M}$ matrix using its singular vectors thus:

$$\mathbf{M} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^{\dagger}. \tag{11}$$

The columns of $\mathbf{V}$ are the right singular vectors (the

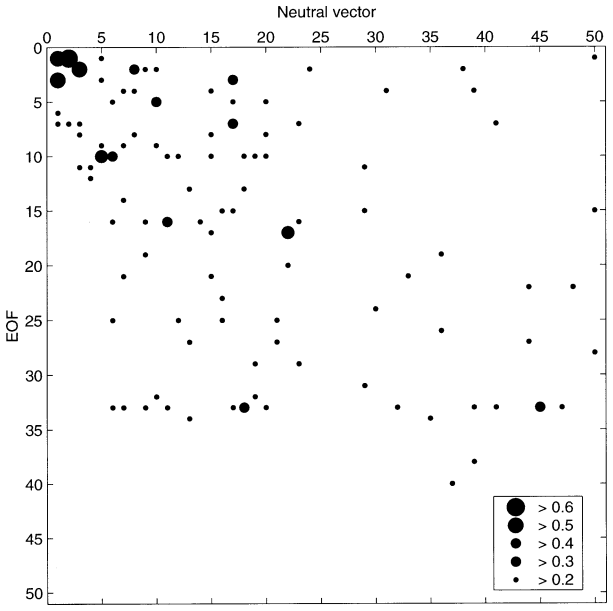

FIG. 6. Cross correlations between model EOFs (vertical axis) and neutral vectors computed using a KE norm (horizontal axis). The absolute value of the correlation is indicated by the size of the dots. Correlations less than 0.2 are not plotted.

neutral vectors); the columns of $\mathbf{U}$ are the left singular vectors (the optimal forcing patterns); and $\mathbf{\Lambda}$ is a diagonal matrix of singular values. Inserting this into (10), and using the fact that $\mathbf{V}^{-1} = \mathbf{V}^{\dagger}$ and $\mathbf{U}^{-1} = \mathbf{U}^{\dagger}$, we obtain

$$\mathbf{\Psi\Psi}^{\dagger} = \mathbf{V}^{-\dagger}\mathbf{\Lambda}^{-1}\mathbf{U}^{-1}\mathbf{U}^{-\dagger}\mathbf{\Lambda}^{-1}\mathbf{V}^{-1} = \mathbf{V\Lambda}^{-2}\mathbf{V}^{\dagger}.$$

Since $\mathbf{V\Lambda}^{-2}\mathbf{V}^{\dagger}$ is in diagonalized form, its eigenvectors are $\mathbf{V}$ and its eigenvalues are the diagonal elements of $\mathbf{\Lambda}^{-2}$. Thus, when $\mathbf{ff}^{\dagger} = \mathbf{I}$, the EOFs (which are the eigenvectors of the covariance matrix $\mathbf{\Psi\Psi}^{\dagger}$) are identical to the singular vectors of the tendency matrix. Since $\mathbf{\Lambda}$ is raised to the -2 power, the dominant EOFs correspond to the smallest singular vectors, that is, to the neutral vectors.

What happens when the transient eddy forcing covariance is *not* proportional to the identity matrix, as we assumed above? Then we have

$$\mathbf{\Psi\Psi}^{\dagger} = (\mathbf{M}^{\dagger}(\mathbf{ff}^{\dagger})^{-1}\mathbf{M})^{-1}. \tag{12}$$

The matrix $(\mathbf{ff}^{\dagger})^{-1}$ is the inverse of the eddy forcing covariance matrix. It is symmetric and positive definite (assuming the inverse exists). It can thus be interpreted as a weight matrix for the inner product between the $\mathbf{M}$ matrix with itself. Thus, in the presence of nonuniform stochastic forcing, we can compare the EOFs to the singular vectors of $\mathbf{M}$ computed using this unusual weight matrix $(\mathbf{ff}^{\dagger})^{-1}$. We demonstrate this by performing an SVD decomposition of $\mathbf{M}$ [Eq. (11)], where now the orthonormality of the $\mathbf{V}$ and $\mathbf{U}$ are defined using the inner products

$$\mathbf{V}^{\dagger}\mathbf{V} = \mathbf{I} \qquad \mathbf{U}^{\dagger}\,(\mathbf{ff}^{\dagger})^{-1}\,\mathbf{U} = \mathbf{I}.$$

Then, using (11):

$$\mathbf{M}^{\dagger}(\mathbf{ff}^{\dagger})^{-1}\mathbf{M} = \mathbf{V\Lambda U}^{\dagger}(\mathbf{ff}^{\dagger})^{-1}\mathbf{U\Lambda V}^{\dagger} = \mathbf{V\Lambda}^{2}\mathbf{V}^{\dagger}.$$

And so, returning to (12):

$$\mathbf{\Psi\Psi}^{\dagger} = (\mathbf{M}^{\dagger}(\mathbf{ff}^{\dagger})^{-1}\mathbf{M})^{-1} = \mathbf{V\Lambda}^{-2}\mathbf{V}.$$

Thus, the principal EOFs of this system are the neutral vectors ($\mathbf{V}$), where the SVD analysis used to compute the neutral vectors employs the inverse of the stochastic forcing correlation matrix as a weight matrix to normalize the optimal forcing patterns ($\mathbf{U}$).

What is the significance of this unusual weight matrix $(\mathbf{ff}^{\dagger})^{-1}$? Consider the simple case where each vector element represents a location in space, and the eddy forcing is spatially uncorrelated but has nonuniform variance. Then the eddy forcing covariance matrix is diagonal, with larger elements on the diagonal where forcing is strong. Thus the weight $(\mathbf{ff}^{\dagger})^{-1}$ is small where the eddy forcing is large. In computing neutral vectors, we want to minimize the forcing needed to excite them. A forcing pattern will be ''small'' with respect to this weighted norm when it has small amplitude where the weight is large, and vice versa. Thus, the SVD analysis selects neutral vectors whose forcing patterns are localized at the site of large eddy forcing.

This allows us to explain and formalize, for example, the observation made in section 3a that the dominant EOFs all lie within the Northern Hemisphere, while the neutral vectors reside in both hemispheres, unless we force them into the north using an artificial hemispheric damping. High-frequency eddy activity, and thus eddy forcing, are far stronger in the Northern Hemisphere (where it is wintertime). Thus, neutral vector patterns sensitive to Northern Hemisphere forcing will be driven more strongly, and so Northern Hemisphere modes will be more prominently visible in the model output, even though they are no more ''neutral'' than Southern Hemisphere modes.

To observe a closer connection between EOFs and neutral vectors, we should compute neutral vectors using a norm weighted with the inverse eddy forcing covariance, rather than any more traditional norm. One approximate way forward would be to use the covariance of high-frequency eddy streamfunction.

4. Optimal forcing patterns

a. Optimal forcing pattern structure

The pattern that maximally excites neutral vector 1 (left column of Fig. 3) shows two broad regions of sensitivity of opposite signs: one focused on Kamchatka and extending over the eastern Pacific and Siberia, and a second focused on the tropical Pacific and covering Africa, Europe, the North Atlantic, and North America. The pattern is baroclinic, with an out-of-phase response in the lower two layers and nearly zero sensitivity at the upper level. This implies a sensitivity to low-level

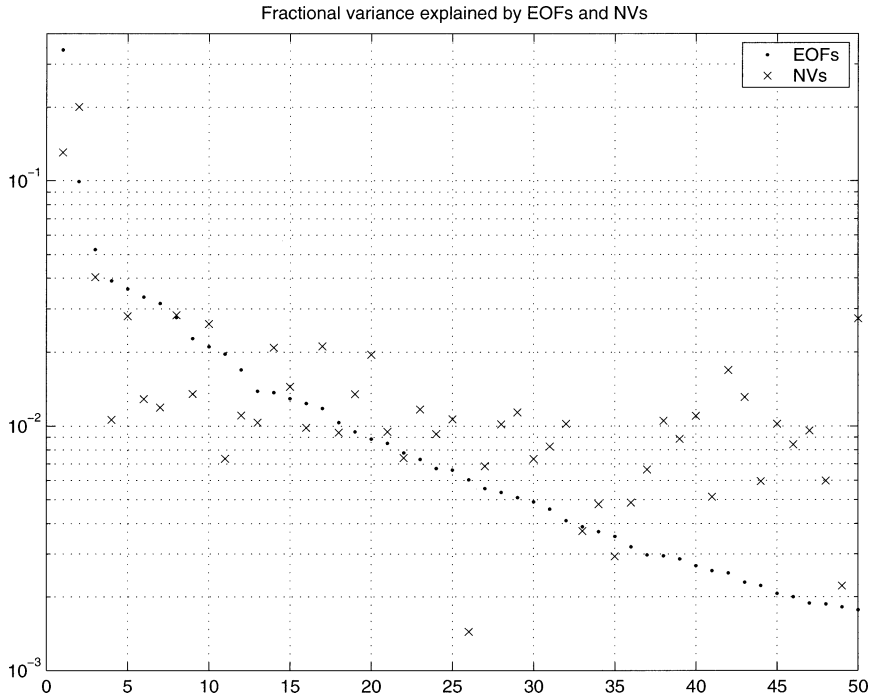

FIG. 7. Comparison of fractional variance of monthly mean model output explained by EOFs (dots) with variance explained by neutral vectors (×).

heating, which hints that this mode may be sensitive to forcing by SST anomalies.

The second optimal forcing pattern (left column of Fig. 4) shows a sensitivity to low level PV forcing over much of the North Pacific, particularly the Sea of Okhotsk, and to an opposite sign of forcing over the North Atlantic and the Arctic. Once again, the mode is sensitive to low-level heating.

The optimal forcing pattern for neutral vector 3 (Fig. 5) also shows a global-scale dipole, with sensitivity to low-level baroclinic forcing. The forcing centers lie at the beginning and end of the PNA-like wave train.

If we stopped here, we would conclude that the optimal forcing patterns are most sensitive to broad-scale, low-level thermal forcing. In some cases (like Fig. 5, in which forcing centers lie at the beginning and end of the wave train), the spatial relationship of forcing to response makes sense, while in others (such as Fig. 3, in which an east–west dipole pattern gives rise to a zonally symmetric structure), it is much less clear.

However, this analysis has chosen to use an inner product in (8) in which the norm of a streamfunction vector is proportional to its kinetic energy. This is only one of a wide array of sensible choices of inner product. For example, we could choose an inner product where the norm of $\boldsymbol{\Psi}$ was proportional to the root-mean-square streamfunction rather than the kinetic energy. We could also attempt to minimize the PV tendency rather than the streamfunction tendency, or use an inner product that applied different weights to different geographical areas or vertical levels.

After considerable experimentation, we find that while the neutral vectors are relatively insensitive to the choice of inner product, the optimal forcing patterns look very different for different inner products. This is demonstrated in Fig. 8, which shows the first neutral vector and optimal forcing pattern using an inner product in which $\|\boldsymbol{\Psi}\|$ is proportional to the root-mean-square streamfunction anomaly, which we call the ''psi norm.''

The first neutral vector for the psi norm looks virtually identical to the first neutral vector of the KE norm. However, the optimal forcing pattern is not the same. Rather than being characterized by a planetary sized dipole structure, we see many narrow, closely spaced zones of positive and negative sensitivity. These tend to be strongest in the Tropics, where, as we have already mentioned, the model's dynamics are the least believable; they also occur to some extent in the Southern Hemisphere (not shown in this polar projection). The tendency for the optimal forcing to be low-level heating remains but is much less prominent. Many of the high and low centers of sensitivity are in similar locations, but generally, the first optimal forcing patterns of the KE and psi norms look quite different.

We have tried several other choices of norm, including an attempt to minimize PV tendency rather than streamfunction tendency; the neutral vectors are not strongly dependent on the norm we choose, while the optimal forcing patterns are highly norm sensitive.

b. Norm sensitivity of optimal forcing patterns

Why do the two norms display identical neutral vectors if their optimal forcing patterns are so different? Consider the atmospheric response equation $\mathbf{M}\Psi = \mathbf{f}$. A neutral vector has small $\mathbf{f}$, so the singular vector decomposition selects the components of Ψ to ensure the near-cancellation of the various terms in the $\mathbf{M}$ matrix. This cancellation is independent of the norm selected, so the neutral vector pattern is not norm sensitive. However, relatively large changes in $\mathbf{f}$ can result from relatively small changes in Ψ, since the left side of the equation is a small difference of large terms. If the norm penalizes one sort of pattern more heavily, the amplitude of that pattern in $\mathbf{f}$ can be made small with only small changes to Ψ. The psi norm applies equal weight to all wavenumbers, while the KE norm penalizes high wavenumbers (for which $\nabla\psi$ is large) more heavily. As a result, the KE optimal forcing patterns are more broad scale. However, this analysis cannot tell us whether one norm is ''better'' than another.

This phenomenon can also be explained in terms of the condition number of the $\mathbf{M}$ matrix (defined as the ratio of the largest and smallest singular values). When this number is large, the matrix is ''poorly conditioned,'' and the response Ψ is sensitive to small perturbations in $\mathbf{f}$ or $\mathbf{M}$. A singular matrix has an infinite condition number. For our $\mathbf{M}$, the condition number is of the order 10^4, implying rather large sensitivity to forcing, but we must stress that the matrix is *not* close to being singular to within machine precision.

Navarra (1993) performed an essentially identical neutral vector analysis on a barotropic model, using the January 300-mb climatological flow. His neutral vector patterns look completely different from those found in this analysis and are much less similar to observed patterns like the NAO, PNA, and AO. However, Navarra's work agrees that the neutral vectors strongly resemble the model's EOFs, and that the condition number of the $\mathbf{M}$ matrix is relatively large, leading to large sensitivity to forcing. Interestingly, Navarra's optimal forcing patterns have some properties similar to those shown in Fig. 8: they are finescale, zonally oriented bands of sensitivity in the Tropics.

c. Response of the nonlinear model to ''optimal'' forcing

We have found in section 4a that the optimal forcing patterns are difficult to define unambiguously, since they are not consistent between different definitions of the inner product. Nevertheless, it is useful to find out whether these patterns, which are optimal in forcing the linearized stationary planetary wave model, are also optimal in forcing the full nonlinear time-evolving model. If this model were quasi-linear, we would expect to see a large neutral vector response to optimal forcing. However, nonlinear effects may come into play. If the optimal forcing patterns found using the linear model do not excite a strong response in the nonlinear model, we cannot expect them to tell us much about the sensitivity of the true atmosphere to PV forcing.

We want to find the time mean perturbation response to a constant ''optimal'' forcing perturbation. We proceed by running three integrations of Molteni's model. First, we create a pair of control runs (runs 1 and 2), where S_n in (5) are unchanged from the specification described in section 2. These two runs are initialized with very slightly different initial conditions and so produce different instances of synoptic eddies. The difference between the mean state of these two runs will give us some idea of the uncertainty of the mean, with which we can compare the experimental run.

In the experimental run (run 3), we perturb S_n by a small amount in the direction of the first optimal forcing pattern for the KE norm (left column of Fig. 3). We use a PV operator to convert the optimal forcing pattern from a streamfunction forcing to a PV forcing. The amplitude of this forcing amounts to about 4% of the basic-state value of S. We expect that the difference in the mean states of runs 1 and 3 should look like KE neutral vector 1 (right column of Fig. 3).

All three runs are performed for 10 000 days of integration; the run length was increased in order to reduce the uncertainty of the sample mean fields.

The top panel of Fig. 9 shows the difference in the means of the two unperturbed runs. We see differences in the means of order 10^5 m^2 s^{-1}. The bottom panel shows the difference between the means of perturbed and unperturbed runs. The differences are likewise of order 10^5 m^2 s^{-1}. Thus, any response seen is indistinguishable from the noise. One might argue that while the amplitude is small, a structure reminiscent of neutral vector 1 is seen in the response. However, this same structure is seen in the top panel too,[3] even though no forcing perturbation was applied there.

The linearized model responds to the forcing with the first neutral vector pattern (right column of Fig. 3), with an amplitude[4] of about 6×10^7 m^2 s^{-1}. The upper limit

[3] Why does the difference between unperturbed runs show neutral vector structure? The uncertainty of the sample mean is $\sigma/\sqrt{N}$, where σ is the standard deviation and N is the number of independent observations. Most atmospheric patterns show little persistence on timescales longer than a week. But the neutral vectors evolve very slowly, persisting for months or longer. Thus the number of *independent* observations of the neutral vectors is abnormally small; the sample mean amplitude of the fast modes converges more quickly than the sample mean amplitude of the neutral vectors as we integrate for longer. In the limit of very long integrations, the residual of sample means will have its structure dominated by the neutral vectors.

[4] This is more than 30% of the amplitude of the model's mean state. While the forcing perturbation is a small fraction of the basic-state forcing, the linear response is a large fraction of the basic-state streamfunction, precisely because the forcing perturbation is optimal. We recognized that a response this large would not fully satisfy linearity conditions; the experiment shown here was a ''first try,'' with an intentionally large forcing to make the response as obvious as possible.

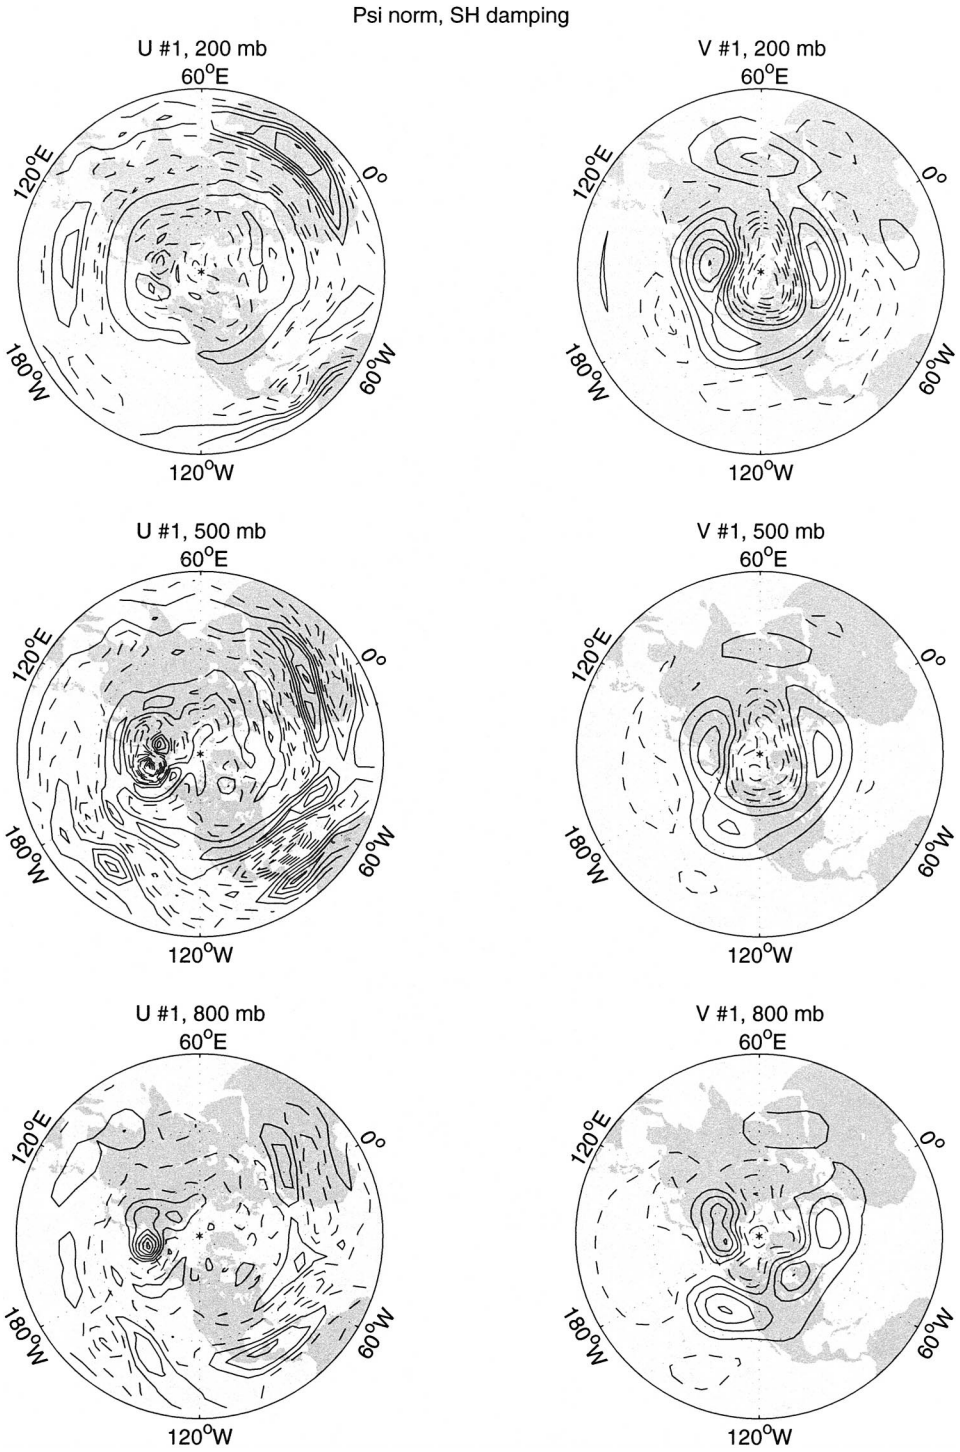

FIG. 8. Right column: Neutral vector 1 for model climatology, psi inner product. Left column: corresponding optimal forcing pattern. Compare with Fig. 3. As in Fig. 3, contour interval is halved for V at 800 mb.

on the nonlinear model's response to the forcing is less than 1% of this. Thus, we conclude that the optimal forcing patterns are ineffective at exciting the nonlinear model.

Why? Take the time average of the PV equation solved by the model (5):

$$J(\overline{\psi}, \overline{q}) = -D(\overline{\psi}) + \overline{S} - \overline{J(\psi', q')}. \qquad (13)$$

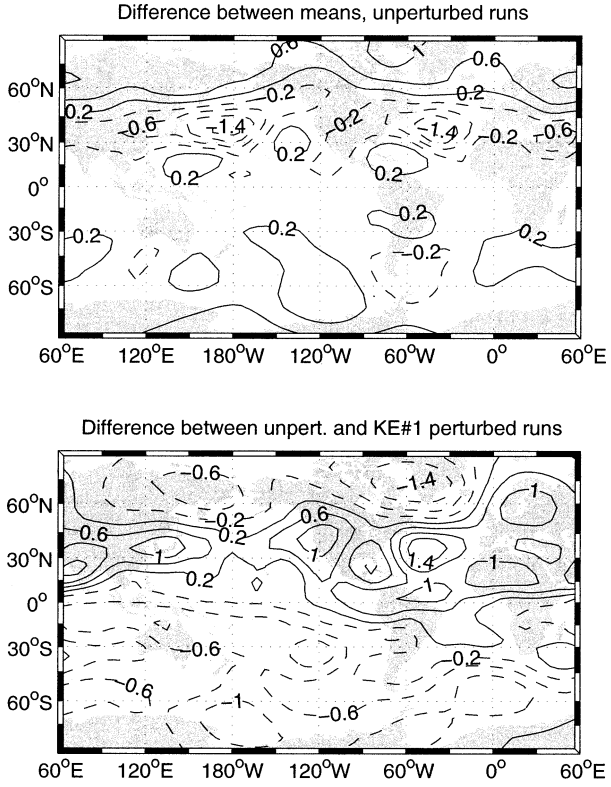

FIG. 9. Response of the nonlinear model to optimal forcing. Top: Streamfunction difference between 10 000-day means of two unperturbed model runs at 500 mb is shown; contour interval is 0.4×10^5 $m^2\ s^{-1}$. Bottom: 500-mb streamfunction response to forcing with the KE optimal forcing pattern 1.

Barred terms represent time means; primed terms represent time-fluctuating terms with zero time mean. Now, consider the time-average balance of PV when we apply a forcing perturbation, S_{nv}. We denote the streamfunction and PV of the response to forcing as ψ_{nv} and q_{nv}:

$$J(\overline{\psi} + \psi_{nv}, \overline{q} + q_{nv}) = -D(\overline{\psi} + \psi_{nv}) + \overline{S} + S_{nv} - \overline{J(\psi', q')}. \quad (14)$$

We expand the terms in (14) and take the difference between it and (13). We assume the perturbation is small enough that the eddy forcing $\overline{J(\psi', q')}$ is unchanged:

$$J(\overline{\psi}, q_{nv}) + J(\psi_{nv}, \overline{q}) + J(\psi_{nv}, q_{nv}) = -D(\psi_{nv}) + S_{nv}.$$

Perform a PV inversion to convert all terms into streamfunction tendencies, as in (6):

$$Q^{-1}[J(\overline{\psi}, q_{nv}) + J(\psi_{nv}, \overline{q}) + D(\psi_{nv})] + Q^{-1}[J(\psi_{nv}, q_{nv})] = f_{nv}.$$

Here, $f_{nv} = Q^{-1}(S_{nv})$ as in (6). The terms that are linear in ψ_{nv} are just the terms that make up the linearized tendency operator **M** in (7):

$$-\mathbf{M}(\psi_{nv}) + Q^{-1}(J(\psi_{nv}, q_{nv})) = f_{nv}. \quad (15)$$

This equation differs from the linear forcing response equation by the presence of the term $Q^{-1}(J(\psi_{nv}, q_{nv}))$. Traditionally, for small forcing perturbations, this term is small. However, we are supplying a forcing that excites a neutral vector. Therefore, both $\mathbf{M}(\psi_{nv})$ and f_{nv} are unusually small. In fact, if we assume the response is linear, and use $-\mathbf{M}(\psi_{nv}) = f_{nv}$ to calculate ψ_{nv} and then check our assumption by plugging that ψ_{nv} into the nonlinear Jacobian term, we find that *the Jacobian term is 600 times larger than the forcing.* The nonlinear dynamics cannot be ignored when forcing is this large. Since the linearized streamfunction response is proportional to the amplitude of the forcing, while the amplitude of the Jacobian term is proportional to the square of the streamfunction, we must make the forcing (and thus the linearized response) 600 times smaller to make $J(\psi_{nv}, q_{nv}) \sim S_{nv}$. This means that the nonlinear self-advection of the response dominates the forcing unless the linearized response amplitude is smaller than 10^5 $m^2\ s^{-1}$. This amplitude is indistinguishable from the sample error of the mean. This means that if we reduce the amplitude of forcing until the response is linear, that response will be too weak to be observed.

A possible solution to this problem was suggested by F. Molteni (2000, personal communication). One could add a quantity to the forcing perturbation S_{nv} that cancels out the contribution of the nonlinear Jacobian term in (15). Thus a large neutral vector response could be excited that did not imbalance the PV equation. One difficulty here is that we carefully selected the optimal forcing pattern to maximally excite a response, but when we add this extra forcing term, the resulting pattern is no longer necessarily optimal in either its linear or nonlinear response. The correct solution to this problem is to perform a *nonlinear* optimization of the response to forcing.

5. Nonlinearity, time dependence, and unstable modes

We have formed our conclusions about the relevance of neutral vectors based on arguments about a system linearized about a mean climate state. In our discussion of the relationships among EOFs, neutral vectors, and optimal forcing patterns, we have often dropped the time-tendency term and considered a steady-state balance, as in (4). But models and the real atmosphere are highly nonlinear systems. They do not settle down into a stable fixed equilibrium with negligible tendency, but continue to fluctuate. Also, the linearized system **M** possesses unstable, rapidly growing modes. From the discussion so far, it is not clear why the linearized, steady-state model should be relevant to systems with these properties. In this section, we discuss a pair of simple systems that share the above properties of atmospheric models and demonstrate that neutral vector/optimal forcing pattern dynamics remains relevant.

A substantial precedent exists for studying the steady linear response of a system using a linearized equation

like (4): this technique was first used to study orographic and thermal forcing of planetary waves by Charney and Eliassen (1949) and Smagorinsky (1953). Models of this sort are able to describe the dominant time mean planetary wave patterns of the atmosphere, even in the presence of nonlinearity, growing modes, and absence of stable fixed equilibria.

However, we need not appeal entirely to precedent. Let us consider the simplest possible system with continual fluctuation about a mean ''climate,'' with an unstable growing mode and nonlinear limitation of the growing mode:

$$\frac{\partial}{\partial t}y = \lambda y - \gamma y|y| + f. \tag{16}$$

Here, λ may be complex, but γ is real and positive. Also, f represents some forcing of the system, which is either constant or much more slowly varying than the internal dynamics.

For $f = 0$, this system has a fixed point at $y = 0$. When $\mathrm{Re}(\lambda) > 0$, this fixed point is unstable, and the system converges onto a limit cycle, making circular orbits about the origin in the complex plane. The time mean of y is zero. When a forcing f is applied, the center of the orbit shifts away from the origin.

We may linearize the model about this time mean, obtaining

$$\frac{\partial}{\partial t}y = \lambda y + f = \mathbf{M}y + f. \tag{17}$$

The above equation is written to resemble (7); the matrix $\mathbf{M}$ is 1×1. In this one-dimensional system, we have only one singular vector, so we cannot test whether neutral vectors are preferred modes of variability; however, we can test the claim that greater neutrality (a smaller singular value of $\mathbf{M}$) leads to increased sensitivity to forcing. The linear system's singular value is simply $\|\lambda\|$; as $\|\lambda\|$ becomes smaller, our earlier discussion claims that response to forcing will increase.

We have performed a large number of forcing response experiments with this system. In each experiment, we choose a different value of λ (we vary the magnitude of λ while keeping the phase angle constant at 45°) and compute the time mean state of the system with and without a constant forcing $f = 0.1$. The difference between the means of the forced and unforced systems is plotted in Fig. 10, as a function of $\|\lambda\|$. As predicted, the response amplitude is largest when $\|\lambda\|$ is small, so that the linearized system is most neutral. Indeed, the amplitude of the response is not far from $f/\|\lambda\|$, which is what a steady-state linear balance [Eq. (17), with $\partial y/\partial t = 0$] would predict.

Why do our linear, steady-state ideas work so well in this nonlinear, fluctuating system? While a linearization is only a good approximation near the point where it was performed, the physical processes it represents (spiralling growth through λ in this example) have an influence over a wide range of phase space. When λ is reduced in this example, the tendency and nonlinear terms must diminish as well, to allow balance. The whole system evolves more sluggishly and weakly. As a result, a forcing f has a greater influence on the system when λ is small.

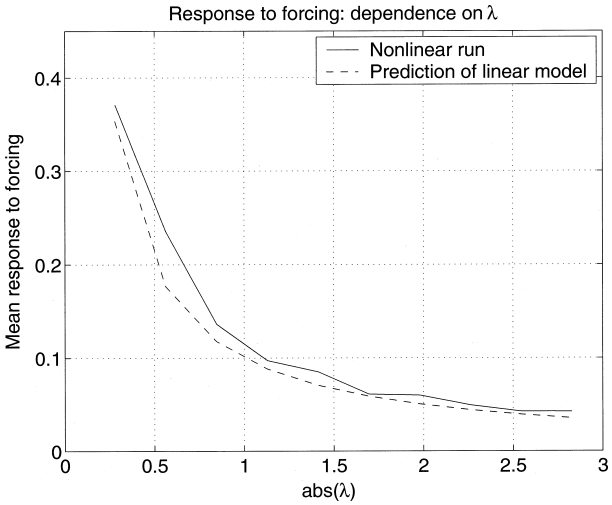

FIG. 10. Amplitude of response of (16) to forcing, as a function of variations in the amplitude of λ. Solid line: amplitude of response to forcing $f = 0.1$ relative to $f = 0$. Dashed: $f/\|\lambda\|$. In all experiments, $\gamma = 1$.

This one-dimensional system is so simple as to be almost trivial, and its stable limit cycle is very dissimilar to the chaotic phase-space trajectories of the atmosphere. We cannot compare neutral vectors to nonneutral patterns in a system with only one singular vector. The Lorenz (1963) system of equations provides a stepping stone toward the full complexity of the atmosphere. Before tackling the three-layer QG model, Marshall and Molteni (1993) performed a neutral vector analysis on the Lorenz system. Linearizing about the time mean of the system, they found that the first neutral vector identifies an axis passing through the two unstable fixed points of the attractor; it is along this line that the Lorenz model exhibits its long-term variability. In addition, Palmer (1999) demonstrated that the maximal response to external forcing in the Lorenz system also lies along this axis. The Lorenz model's long-term variability and response to forcing are complicated and fundamentally nonlinear; nevertheless, the neutral vector accurately identifies their patterns. This increases confidence in neutral vector analysis as a tool for analyzing nonlinear systems.

6. Conclusions

1) Is there a connection between the leading modes of variability of the model and the neutral modes of a linearized version of the model?

Yes. The neutral vectors of Molteni's three-layer model closely resemble the EOFs of that model, even though the EOFs are computed from data generated by a nonlinear model integration and the neutral vectors arise from singular vector decomposition of a linearized model. The neutral vectors explain almost as much of the nonlinear model's variance as the EOFs do; this indicates their importance in understanding model variability and makes them useful in constructing reduced-subspace models of atmospheric variability.

We have computed our neutral vectors using a model linearized using the model's time mean as a basic state. However, the time mean is a statistical construct: we will never actually observe the atmosphere in this state, and it may not even lie on the attractor of this chaotic system. Thus, it is rather surprising that we get such a strong similarity between our neutral vectors and the model's EOFs. However, other reasonable choices of basic state (for instance, a snapshot of the nonlinear model state at a single time point) produce neutral vectors with no similarity to EOFs of the model or observations.

2) Why should such a connection exist?

Neutral vector patterns exhibit a maximal response to forcing perturbations. If one assumes that the EOFs result from linear excitation of relatively slow modes by transient eddy PV forcing, then one can demonstrate that the patterns of EOFs and neutral vectors are mathematically identical. This is an exact identity only if the neutral vectors are computed using an inner product weighted by the inverse of the eddy forcing covariance, but remains approximately true for other inner products.

3) Do the model's neutral modes resemble observed modes of strong (and possibly coupled) variability?

Yes. We find neutral-vector analogues to the PNA, AO, and (to a lesser extent) NAO patterns.

Since neutral vectors resemble EOFs of the model and of observed fields, they are a very useful concept in studying the low-frequency variability of the atmosphere. Their advantage over EOFs is that they have a simple physical interpretation: they are the most slowly evolving patterns within the atmosphere, the patterns for which advection of PV anomalies nearly balances their tendency to propagate as Rossby waves.

These results lead to the following tempting, though speculative, chain of logic: Fluctuations of EOF patterns like the NAO and PNA account for much of the atmosphere's interannual variability. The EOFs of the observed atmosphere look like this model's neutral vectors. Since the neutral vectors are, in the linear model, the patterns that most strongly respond to forcing, we are led to suspect that if any atmospheric modes are involved in interannual atmosphere–ocean coupled interactions, the NAO and its relatives should be.

We have previously discussed a simple model of interannual coupled variability (Goodman and Marshall 1999; GM02). The prominent coupled modes generated within that model obey neutral vector dynamics in the model atmosphere, in accordance with the argument above. However, even if the specific coupled dynamics described in (Goodman and Marshall 1999; GM02) are not active in the real world, the high linear sensitivity of the neutral vectors makes them a useful paradigm for the study of atmosphere–ocean interaction; this is discussed in detail in GM02.

Observations (Walker and Bliss 1932; Hurrell 1995) show that the same patterns of variability dominate both the intraseasonal and the interannual variability of the atmosphere. The neutral vector paradigm gives one of several possible explanations for this. Neutral vectors are the patterns that respond most strongly to forcing in the linear model. High-frequency forcing from transient eddies will generate short-term variability, while low-frequency forcing from SST anomalies or other processes will lead to interannual variability. The same pattern is easily excited by both intrinsic and extrinsic forcing of the atmosphere. Hurrell (1996) finds that much of the recent Northern Hemisphere warming can be explained by a secular change in the amplitude of naturally occurring patterns like the NAO and PNA. Some interpret this as evidence that global warming is simply random variation in a natural phenomenon, but the present work suggests that if these patterns are preferred modes of variability, they may be easily excited by anthropogenic forcing. Intraseasonal variability and long-term climate change may have the same distinctive patterns, but be excited by different processes.

4) What forcing patterns optimally excite these modes?

While neutral vectors are a powerful tool for understanding the atmosphere's variability, the optimal forcing patterns appear to be much less useful. While their structure suggests that the neutral vectors may be easily excited by low-level baroclinic forcing, such as would arise from heating generated by SST anomalies, their patterns are not robust. We can get very different optimal forcing patterns by using a different inner product in our singular vector decomposition. This is because the optimal forcing patterns are the small residual differences between large terms in the atmospheric response equation, so small differences in the analysis can make huge differences in the results. The neutral vectors do not display this sensitivity. Even worse, the optimal forcing patterns are ineffective in exciting a large response in the full nonlinear model. This is because the nonlinear self-interaction of the response to forcing dominates over the forcing itself, unless that response is tiny.

This may make it difficult to generalize the coupled atmosphere–ocean interaction that was discussed in Goodman and Marshall (1999) and GM02 to more complicated models of air–sea interaction. While the neutral vectors are the most sensitive to forcing (including forc-

ing arising from SST anomalies) in linear models, they are not very responsive in the nonlinear model discussed in section 4c. In addition, the discussion of the coupled interaction between neutral vectors and the ocean in GM02 hinged on the idea that, as SST anomalies moved around, they projected onto first one, then another of the optimal forcing patterns. Each strong projection excited a neutral vector response, which fed back onto the ocean. Here, we find that the shape of the optimal forcing patterns is not very robust. As a result, the regular alternation of projection necessary for the coupled wave may be difficult to identify, if it occurs at all.

Acknowledgments. We would like to thank Franco Molteni, Fabio d'Andrea, and two anonymous reviewers for helpful advice. David Ferreira provided us with model code. This work was supported by the NOAA Office of Global Programs.

REFERENCES

Achatz, U., and G. Branstator, 1999: A two-layer model with empirical linear corrections and reduced order for studies of internal climate variability. *J. Atmos. Sci.,* **56,** 3140–3160.

Andersson, P., M. Berggren, and D. S. Henningson, 1999: Optimal disturbances and bypass transition in boundary layers. *Phys. Fluids,* **11** (1), 134–150.

Barnston, A., and R. E. Livezey, 1987: Classification, seasonality and persistence of low-frequency circulation patterns. *Mon. Wea. Rev.,* **115,** 1083–1126.

Cayan, D., 1992: Latent and sensible heat flux anomalies over the northern oceans: The connection to monthly atmospheric circulation. *J. Climate,* **5,** 354–369.

Charney, J. G., and A. Eliassen, 1949: A numerical method for predicting the perturbations of middle latitude westerlies. *Tellus,* **1,** 38–54.

Czaja, A., and J. Marshall, 2001: Observations of atmosphere–ocean coupling in the North Atlantic. *Quart. J. Roy. Meteor. Soc.,* **127,** 1893–1916.

d'Andrea, F., 2002: Extratropical low-frequency variability as a low dimensional problem. Part II: Stationarity and stability of large scale equilibria. *Quart. J. Roy. Meteor. Soc.,* **128,** 1059–1073.

Farrell, B. F., 1982: The initial growth of disturbances in a baroclinic flow. *J. Atmos. Sci.,* **39,** 1663–1686.

——, 1989: Optimal excitation of baroclinic waves. *J. Atmos. Sci.,* **46,** 1193–1206.

Goodman, J. C., 2001: Interannual middle-latitude atmosphere–ocean interactions. Ph.D. thesis, Massachusetts Institute of Technology, 151 pp.

——, and J. Marshall, 1999: A model of decadal middle-latitude atmosphere–ocean coupled modes. *J. Climate,* **12,** 621–641.

——, and ——, 2002: The role of neutral singular vectors in middle-latitude air–sea coupling. *J. Climate,* in press.

Haines, K., and A. Hannachi, 1995: Weather regimes in the Pacific from a GCM. *J. Atmos. Sci.,* **52,** 2444–2462.

Hannachi, A., 1997: Low-frequency variability in a GCM: Three-dimensional flow regimes and their dynamics. *J. Climate,* **10,** 1357–1379.

Hasselmann, K., 1988: PIPs and POPs: The reduction of complex dynamical systems using principal interaction and oscillation patterns. *J. Geophys. Res.,* **83,** 11 015–11 021.

Hurrell, J. W., 1995: Decadal trends in the North Atlantic Oscillation: Regional temperatures and precipitation. *Science,* **269,** 676–679.

——, 1996: Influence of variations in extratropical wintertime teleconnections on Northern Hemisphere temperature. *Geophys. Res. Lett.,* **23** (6), 665–668.

Kaplan, A., Y. Kushnir, and M. A. Cane, 2000: Reduced space optimal interpolation of historical marine sea level pressure: 1854–1992. *J. Climate,* **13,** 2987–3002.

Kushnir, Y., 1994: Interdecadal variations in North Atlantic sea surface temperature and associated atmospheric conditions. *J. Climate,* **7,** 141–157.

Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.,* **20,** 130–141.

Luchini, P., 2000: Reynolds-number-independent instability of the boundary layer over a flat surface: Optimal perturbations. *J. Fluid Mech.,* **404,** 289–309.

Marshall, J., and F. Molteni, 1993: Toward a dynamical understanding of planetary-scale flow regimes. *J. Atmos. Sci.,* **50,** 1792–1818.

Michelangeli, P.-R., and R. Vautard, 1998: The dynamics of Euro-Atlantic blocking onsets. *Quart. J. Roy. Meteor. Soc.,* **124,** 1045–1070.

Molteni, F., 1994: Towards a dynamical understanding of planetary-scale flow regimes. Ph.D. thesis, Imperial College, University of London, 196 pp.

——, and T. Palmer, 1993: Predictability and finite-time instability of the northern winter circulation. *Quart. J. Roy. Meteor. Soc.,* **119,** 269–298.

——, A. Sutera, and N. Tronci, 1988: The EOFs of the geopotential eddies at 500 mb in winter and their probability density distributions. *J. Atmos. Sci.,* **45,** 3063–3080.

Navarra, A., 1993: A new set of orthonormal modes for linearized meteorological problems. *J. Atmos. Sci.,* **50,** 2569–2583.

North, G. R., 1984: Empirical orthogonal functions and normal modes. *J. Atmos. Sci.,* **41,** 879–887.

Palmer, T. N., 1999: A nonlinear dynamical perspective on climate prediction. *J. Climate,* **12,** 575–591.

Richman, M. B., 1986: Rotation of principal components. *J. Climatol.,* **6** (3), 293–335.

Smagorinsky, J., 1953: The dynamical influence of large-scale heat sources and sinks on the quasi-stationary mean motion of the atmosphere. *Quart. J. Roy. Meteor. Soc.,* **97,** 342–366.

Thompson, D. W. J., and J. M. Wallace, 1998: The Arctic Oscillation signature in the wintertime geopotential height and temperature fields. *Geophys. Res. Lett.,* **25** (9), 1297–1300.

Walker, G., and E. Bliss, 1932: World weather V. *Mem. Roy. Meteor. Soc.,* **4,** 53–83.

Wallace, J. M., 2000: North Atlantic Oscillation/annular mode: Two paradigms—one phenomenon. *Quart. J. Roy. Meteor. Soc.,* **126,** 791–805.

——, Y. Zhang, and K. H. Lau, 1993: Structure and seasonality of interannual and interdecadal variability of the geopotential height and temperature fields in the Northern Hemisphere troposphere. *J. Climate,* **6,** 2063–2082.

Whitaker, J. S., and P. D. Sardeshmukh, 1998: A linear theory of extratropical synoptic eddy statistics. *J. Atmos. Sci.,* **55,** 237–258.

White, W. B., and R. G. Peterson, 1996: An Antarctic circumpolar wave in surface pressure, wind, temperature and sea-ice extent. *Nature,* **380,** 699–702.