Chapter 3

Statistical Wind Stress Model

A field of water betrays the spirit that is in the air.

Thoreau, *Walden*

3.1 Motivation

The simulation and prediction of the El Niño–Southern Oscillation (ENSO) remains a formidable challenge (Landsea and Knaff, 2000; Latif et al., 2001). Part of the challenge stems from uncertainties about the tropical Pacific wind stress response to sea surface temperature anomalies (SSTA). Several studies have shown that this response is critical to the behavior of ENSO (Neelin, 1990; Kirtman, 1997; An and Wang, 2000; Cassou and Perigaud, 2000). Unfortunately, general circulation models of the atmosphere have had difficulty reproducing the observed response (Saji and Goswami, 1997; Kleeman et al., 2001), and even the observational analyses themselves do not fully agree on what reality is like (Chapter 2; McPhaden et al., 1998).

To solve this problem, several researchers have constructed statistical models of the wind stress based on the observational analyses. These statistical atmospheres may then be coupled to dynamical ocean models to produce "hybrid coupled models" (Barnett et al., 1993; Balmaseda et al., 1994; Syu and Neelin, 1995; Blanke et al., 1997; Eckert and Latif, 1997; Cassou and Perigaud, 2000; Kang and Kug, 2000). Recently, Harrison et al. (2002) built several such hybrid models, coupling statistical atmospheres derived from various flux products to the same ocean GCM. They found that different atmospheric responses produced very different ENSO behaviors, ranging from damped oscillations to sustained oscillations to a slow drift to a new steady state. The predictive skill of these coupled models varied as well. Such sensitivity motivated the careful investigation of the differences between available tropical Pacific wind stress products in Chapter 2, and also warrants a closer look at the statistical models derived from them.

A statistical model is much more than a "black box"— it is a powerful diagnostic tool, as it gathers a researcher's qualitative notions about a system into a consistent, quantitative framework. This tool can then be used to test ideas, e.g. to reveal the influence of SST

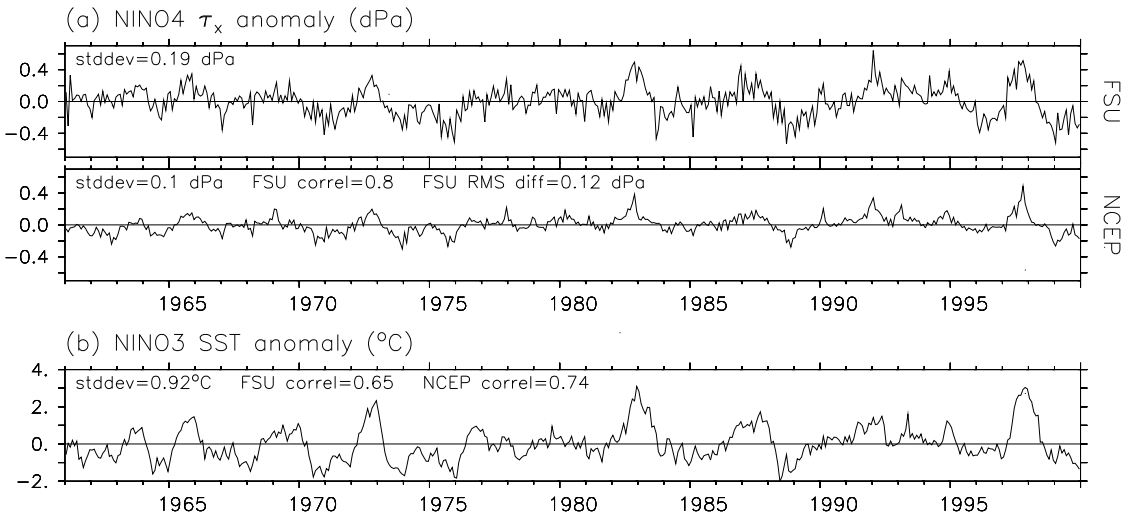

Figure 3.1: Timeseries of monthly-mean (a) zonal wind stress anomalies averaged over the NINO4 region (160°E–150°W, 5°S–5°N) and (b) SST anomalies averaged over the NINO3 region (150°W–90°W, 5°S–5°N). Anomalies are with respect to the 1961–1999 climatology. The standard deviation of each timeseries is indicated, followed by correlations and/or root-mean-square differences with the timeseries in the other panels.

forcing, nonlinearity, and noise on the observed wind stress, even (or perhaps especially) when the model is found to be incomplete. Model-building and evaluation thus represent the essence of the scientific method.

This study therefore derives statistical models from two different widely-used analyses. We assume the monthly-mean wind stress consists of three independent components: a seasonally-varying climatology, a linear response to SSTA, and random noise. After fitting this statistical model to the analyses for two different time periods, we investigate to what extent such a model actually reproduces the data. Particular attention is paid to the sizable residual that remains after SSTA dependence is subtracted from the stress anomalies, and prior hypotheses about this residual are tested.

3.2 Data

The data used in this study are the same as in the previous chapter. Fig. 3.1 shows timeseries of area-averaged monthly τ_x' and SSTA with respect to the 1961–1999 base period. The area averages comprise the regions of maximum variability for τ_x' (NINO4: 160°E–150°W, 5°S–5°N) and SSTA (NINO3: 150°W–90°W, 5°S–5°N). For both the FSU and NCEP wind stress products, τ_x' is well-correlated with SSTA, and ENSO events are stronger and have a longer period toward the latter half of the record. What is striking is that the RMS difference between the FSU and NCEP anomalies is of the same order as the stress anomalies themselves, and the correlation of FSU and NCEP with each other is not much larger than their correlation with SSTA. The FSU stress anomalies are stronger,

noisier, less correlated with SSTA, and have more persistent westerly peaks than NCEP. Prominent differences between the products are evident at the beginning of the record (1961–1965), when FSU indicates westerlies while NCEP shows easterlies. Differences are also apparent at the end of the record, particularly during the winter of 1991–1992 when FSU shows stronger westerly anomalies than NCEP, and in 1996 when FSU shows stronger easterly anomalies than NCEP. A more complete account of the differences between the FSU and NCEP products may be found in Chapter 2.

Several things are evident from Fig. 3.1. First, it is clear that subannual variance is an important component of $\boldsymbol{\tau}'$. Second, this subannual variability is probably not related to the large-scale monthly-mean SSTAs, which vary mostly on interannual time scales. Lastly, differences in $\boldsymbol{\tau}'$ between the analyses are as apparent as changes between decades. This suggests any "real" changes in the stress anomalies will be difficult to extract from the "observed" changes, since the former are confounded by changes in the observing system and analysis methods. In this study, we focus not on the stress anomalies themselves, but instead on their relationship with SSTA. We hope to determine whether this relationship differs between the observational analyses, and whether there has been any change in this relationship in recent years.

As discussed in Chapter 2, most of the interannual variability of the wind stress is at key climatological "edges" where strong SST gradients exist; the warm pool/cold tongue boundary and the southern edge of the ITCZ are two such regions. Local nonlinearity may be important when these edges shift positions during ENSO events. It is therefore necessary to assess the adequacy of commonly-used linear models of the wind stress. To this end, we construct such a model and assess its validity.

3.3 Empirical bases for climate variability

This section reviews common methods for building empirical bases for observed climate variability, and then applies these methods separately to each data product and time period. Singular value decomposition (SVD) will provide a convenient basis for the covariability of SST and wind stress, while principal components analysis (PCA) will provide a basis for the part of the stress not linearly coupled to SST.

3.3.1 Singular value decomposition

Towards the development of the deterministic model, we briefly describe SVD and its use for finding coupled patterns in data. Further details are given by Bretherton et al. (1992), Newman and Sardeshmukh (1995), and Cherry (1996, 1997).

Let $\mathbf{X}$ be an $n \times p$ data matrix, whose rows represent snapshots of a data field at different times. The data values are denoted x_{li} where $l = 1, \ldots, n$ is the time index and $i = 1, \ldots, p$ is the space index. We assume $\mathbf{X}$ is centered in time, namely $\sum_{l=1}^{n} x_{li}/n = 0$ for all spatial gridpoints i. Let $\mathbf{Y}$ be another time-centered data matrix with dimensions $n \times q$, and let $'$ denote a matrix transpose. Then the cross-covariance of $\mathbf{X}$ and $\mathbf{Y}$ is given by the $p \times q$ matrix

$$\mathbf{C} \equiv \frac{\mathbf{X}'\mathbf{Y}}{n-1} \tag{3.1}$$

By definition, SVD produces

$$\mathbf{C} = \widetilde{\mathbf{A}}\mathbf{D}\widetilde{\mathbf{B}}' \tag{3.2}$$

where nondimensional matrices are tagged with a tilde ($\sim$). $\mathbf{D}$ is a diagonal matrix of dimension $r \times r$, where $r \equiv \min(p, q)$. The diagonal elements of $\mathbf{D}$ are called the *singular values* of $\mathbf{C}$, denoted d_k for $k = 1, \ldots, r$. The matrices $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ have dimensions $p \times r$ and $q \times r$ and are unitary:

$$\widetilde{\mathbf{A}}'\widetilde{\mathbf{A}} = \widetilde{\mathbf{B}}'\widetilde{\mathbf{B}} = \mathbf{I} \tag{3.3}$$

where $\mathbf{I}$ is the $r \times r$ identity matrix. The columns of $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{B}}$ are called the left and right *singular vectors* of $\mathbf{C}$, denoted $\widetilde{\mathbf{a}}_k$ and $\widetilde{\mathbf{b}}_k$.

Timeseries of singular vector expansion coefficients are given by the $n \times r$ matrices

$$\mathbf{X}^* \equiv \mathbf{X}\widetilde{\mathbf{A}} \tag{3.4}$$

$$\mathbf{Y}^* \equiv \mathbf{Y}\widetilde{\mathbf{B}} \tag{3.5}$$

Thus the first column of $\mathbf{X}^*$, call it $\mathbf{x}_1^*$, is the timeseries resulting from projection of the data $\mathbf{X}$ onto the first left singular vector $\widetilde{\mathbf{a}}_1$. Using the expansion timeseries, one constructs covariance matrices

$$\mathbf{C}_{x^*x^*} \equiv \frac{\mathbf{X}^{*\prime}\mathbf{X}^*}{n-1} \tag{3.6}$$

$$\mathbf{C}_{y^*y^*} \equiv \frac{\mathbf{Y}^{*\prime}\mathbf{Y}^*}{n-1} \tag{3.7}$$

$$\mathbf{C}_{x^*y^*} \equiv \frac{\mathbf{X}^{*\prime}\mathbf{Y}^*}{n-1} = \mathbf{D} \tag{3.8}$$

The last equality, which results from application of (3.1)–(3.5), shows that d_k is the covariance of the expansion timeseries $\mathbf{x}_k^*$ and $\mathbf{y}_k^*$, and that $\mathbf{x}_{k_1}^*$ and $\mathbf{y}_{k_2}^*$ are uncorrelated for $k_1 \neq k_2$.

One may define a "length" for $\mathbf{C}$ using the Frobenius norm

$$\|\mathbf{C}\|^2 \equiv \sum_{i=1}^{p}\sum_{j=1}^{q} c_{ij}^2 = \sum_{k=1}^{r} d_k^2 \tag{3.9}$$

where the second equality, shown by Strang (1988), is simply an application of the Pythagorean theorem. Now suppose one constructs synthetic timeseries $\mathbf{X}_N$ and $\mathbf{Y}_N$, using only the leading N pairs of singular vectors. Call the covariance matrix of these synthetic timeseries $\mathbf{C}_N$. The cumulative fraction of squared covariance accounted for the first N modes is then given by

$$\frac{\|\mathbf{C}_N\|^2}{\|\mathbf{C}\|^2} = \frac{\sum_{k=1}^{N} d_k^2}{\sum_{k=1}^{r} d_k^2} \tag{3.10}$$

It has been shown (Stewart, 1973) that SVD is in fact the *optimal* way to account for $\|\mathbf{C}\|^2$ with N pairs of patterns. In other words, the leading singular mode patterns $\mathbf{a}_1$ and $\mathbf{b}_1$ have expansion timeseries with the maximum possible covariance of any pair of patterns $\mathbf{a}$ and $\mathbf{b}$, and each successive mode has the maximum possible covariance of any

pair of patterns orthogonal to all earlier pairs. In this sense, SVD of the covariance matrix produces an efficient set of empirical "coupled modes." There is no reason to expect these modes to represent actual *physical* relationships between the data fields, except to the extent that these relationships impart strongly covariant structures in the data (Newman and Sardeshmukh, 1995).

For convenience, we transfer the physical units from the singular vector expansion timeseries to the singular vector patterns. It is convenient to define

$$\begin{aligned} \mathbf{S}_{x^*}^2 &\equiv \text{diag}\,(\mathbf{C}_{x^*x^*}) \\ \mathbf{S}_{y^*}^2 &\equiv \text{diag}\,(\mathbf{C}_{y^*y^*}) \end{aligned}$$

where the "diag" operator simply sets off-diagonal elements to zero. Pre(post)-multiplication of a matrix by $\mathbf{S}_{x^*}$ then corresponds to multiplying the k^{th} row (column) by the standard deviation of the k^{th} expansion coefficient. Dimensional singular vector matrices $\mathbf{A}$ and $\mathbf{B}$, and nondimensional expansion coefficient matrices $\widetilde{\mathbf{X}}^*$ and $\widetilde{\mathbf{Y}}^*$, are then constructed as follows:

$$\mathbf{A} \equiv \widetilde{\mathbf{A}}\mathbf{S}_{x^*} \tag{3.11}$$

$$\mathbf{B} \equiv \widetilde{\mathbf{B}}\mathbf{S}_{y^*} \tag{3.12}$$

$$\widetilde{\mathbf{X}}^* \equiv \mathbf{X}^*\mathbf{S}_{x^*}^{-1} \tag{3.13}$$

$$\widetilde{\mathbf{Y}}^* \equiv \mathbf{Y}^*\mathbf{S}_{y^*}^{-1} \tag{3.14}$$

By construction, the nondimensional expansion coefficients $\widetilde{\mathbf{x}}_k^*$ and $\widetilde{\mathbf{y}}_k^*$ have unit variance, and cross-correlation

$$\widetilde{\mathbf{D}} \equiv \mathbf{S}_{x^*}^{-1}\mathbf{D}\mathbf{S}_{y^*}^{-1} \tag{3.15}$$

With these definitions, the singular value decomposition (3.2) takes the form

$$\mathbf{C} = \mathbf{A}\widetilde{\mathbf{D}}\mathbf{B}' \tag{3.16}$$

In summary, SVD efficiently partitions the covariance of $\mathbf{X}$ and $\mathbf{Y}$ into a particular set of patterns, namely the r columns of $\mathbf{A}$ and $\mathbf{B}$. The first column of $\mathbf{A}$ is paired with the first column of $\mathbf{B}$, the second with the second, and so on, with successive pairs exhibiting successively less temporal covariance between their respective expansion timeseries. The correlations of these paired timeseries are listed along the diagonal of $\widetilde{\mathbf{D}}$. All other pairings of $\mathbf{A}$ expansion timeseries with $\mathbf{B}$ expansion timeseries are uncorrelated. The patterns of $\mathbf{A}$ are mutually orthogonal but their expansion timeseries may be correlated; the same goes for $\mathbf{B}$.

3.3.2 Principal components analysis

PCA, also called empirical orthogonal function (EOF) analysis, is simply SVD of the data matrix $\mathbf{X}$:

$$\mathbf{X} = \widetilde{\mathbf{A}}_x\mathbf{D}_x\widetilde{\mathbf{B}}_x' \tag{3.17}$$

$$\widetilde{\mathbf{A}}_x'\widetilde{\mathbf{A}}_x = \widetilde{\mathbf{B}}_x'\widetilde{\mathbf{B}}_x = \mathbf{I} \tag{3.18}$$

$\mathbf{D}_x$ is a diagonal matrix of dimension $s \times s$, where $s \equiv \min(n, p)$. $\widetilde{\mathbf{A}}_x$ and $\widetilde{\mathbf{B}}_x$ have dimensions $n \times s$ and $p \times s$. The columns of $\widetilde{\mathbf{B}}_x$ are orthonormal patterns, with expansion coefficients

$$\mathbf{X}_x^* \equiv \widetilde{\mathbf{A}}_x \mathbf{D}_x \tag{3.19}$$

From (3.18), it follows that the cross-covariance matrix of these expansion coefficients is

$$\frac{\mathbf{X}_x^{*\prime} \mathbf{X}_x^*}{n-1} = \frac{\mathbf{D}_x^2}{n-1} \tag{3.20}$$

and so the PCA expansion coefficients are uncorrelated. The expansion coefficient variances, given by the diagonal elements of $\mathbf{D}_x^2/(n-1)$, are sequentially maximized and sum to the total variance $\|\mathbf{X}\|^2/(n-1)$.

For convenience, we transfer the physical units from the expansion timeseries to the patterns using

$$\mathbf{B}_x \equiv (n-1)^{-1/2} \widetilde{\mathbf{B}}_x \mathbf{D}_x \tag{3.21}$$

$$\widetilde{\mathbf{X}}_x^* \equiv (n-1)^{1/2} \mathbf{X}_x^* \mathbf{D}_x^{-1} = (n-1)^{1/2} \widetilde{\mathbf{A}}_x \tag{3.22}$$

The nondimensional expansion coefficients so defined have unit variance. With these definitions, the principal components decomposition (3.17) takes the form

$$\mathbf{X} = \widetilde{\mathbf{X}}_x^* \mathbf{B}_x' \tag{3.23}$$

In summary, PCA efficiently partitions the variance of $\mathbf{X}$ into a particular set of patterns, namely the s columns of $\mathbf{B}_x$, with successive patterns exhibiting successively less temporal variance of their expansion timeseries. These patterns are mutually orthogonal and their expansion timeseries are uncorrelated.

3.3.3 SVD applied to the data

In this section we apply SVD to the data, following Gutzler (1993) and Syu and Neelin (1995). The singular vectors of SST will serve as a convenient set of basis vectors for the deterministic wind stress model of Section 3.4.

The anomalies of SST and wind stress are arranged as the rows of $\mathbf{X}$ and $\mathbf{Y}$ as in the last section. The first half of each row of $\mathbf{Y}$ contains τ_x' and the second half contains τ_y'. The 1961–1979 and 1980–1999 periods and the NCEP and FSU wind stress products are treated separately, giving four different cases. The time dimension is $n = 228$ for 1961–1979, and $n = 240$ for 1980–1999. The spatial matrix dimensions are $p = 494$, and $q = 494 \times 2 = 988$ for all cases.

The cumulative fraction of square covariance accounted for by the singular modes is shown in Fig. 3.2. The singular modes efficiently capture the covariation of SST and wind stress regardless of the dataset or time period, with the first 8 modes accounting for over 99% of the square covariance. That the gravest mode accounts for over 84% of the square covariance demonstrates that to first order, ENSO is a standing oscillation. Interestingly, the modes account for the FSU/SSTA covariance more easily than they account for the NCEP/SSTA covariance.

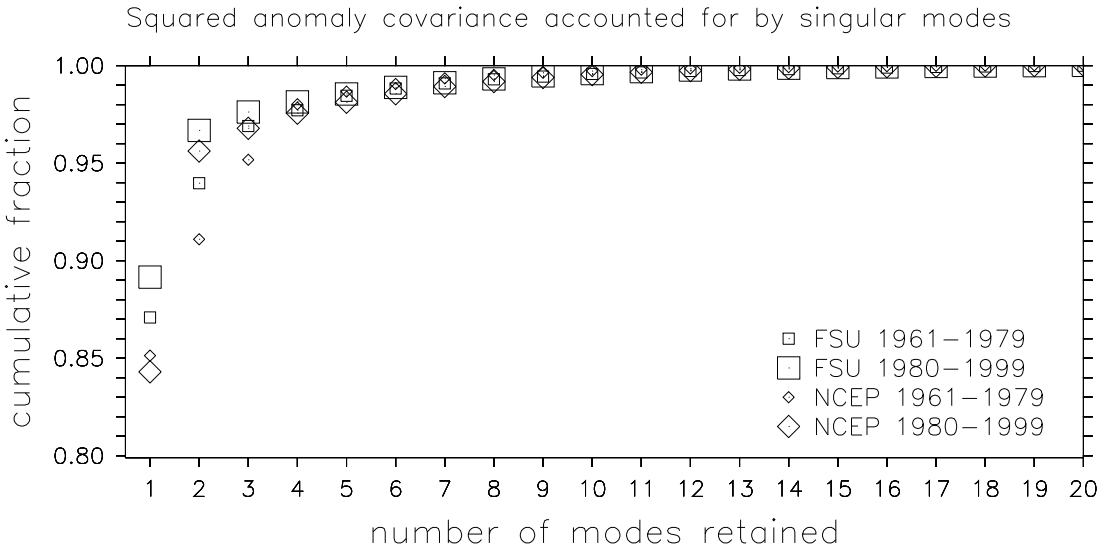

Figure 3.2: Cumulative fraction of squared covariance of monthly-mean tropical Pacific SST and wind stress anomalies, accounted for by the first 20 singular modes.

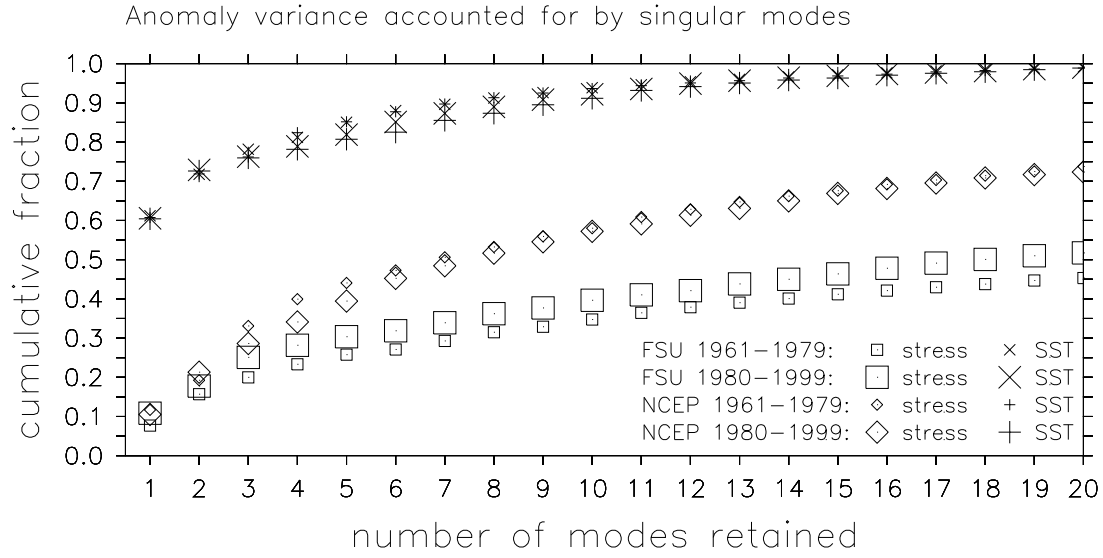

Figure 3.3: Cumulative fraction of variance of monthly-mean tropical Pacific SST and wind stress anomalies, accounted for by the first 20 singular modes.

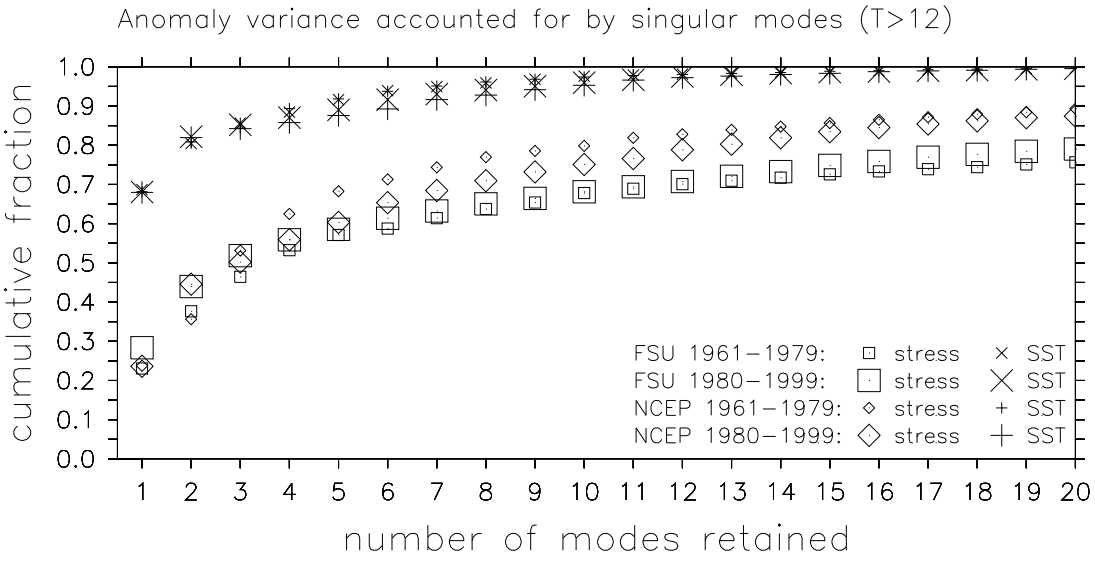

Figure 3.4: Cumulative fraction of interannual variance of tropical Pacific SST and wind stress anomalies, accounted for by the first 20 singular modes.

The singular modes are less efficient at capturing the variances of the fields (Fig. 3.3). The first 8 modes account for more than 85% of the SST variance, but less than 35% of the stress variance for FSU and less than 50% of that for NCEP. This is partly because most of the stress variance occurs at subannual periods irrelevant to SST changes. When the timeseries are filtered to retain only periods greater than one year, the stress variance captured by the first 8 modes increases substantially (Fig. 3.4). However, even on interannual time scales a large fraction of the wind stress does not covary with SSTA. With regard to the explained variance, there is a bigger difference between the FSU and NCEP products than between the 1961–1979 and 1980–1999 time periods, with the SVD capturing less of the total variance in FSU. This is related to the "noisiness" of the FSU stress, which was evident in Fig. 3.1.

The singular vectors $\mathbf{a}_1$, $\mathbf{b}_1$ and $\mathbf{a}_2$, $\mathbf{b}_2$ are shown for each dataset and period in Fig. 3.5. The expansion timeseries for mode 1 (not shown) track standard ENSO indices, such as the NINO3 SSTA and NINO4 τ_x' of Fig. 3.1. In its positive phase, mode 1 is reminiscent of an El Niño event. There is equatorial warming in eastern Pacific, with westerly and equatorward wind anomalies in the central Pacific which converge onto the warm SSTA. Weak easterly anomalies are present over the region of strongest warming in the east, but there are almost no zonal stress anomalies near the coasts.

The SST and wind stress patterns of mode 2 are roughly in quadrature with those of mode 1. Therefore any large-scale propagation of anomalies will be evident in the temporal phasing of the mode expansion coefficients. Fig. 3.6 shows the correlation between the expansion timeseries of mode 1 and those of lagged mode 2. Interestingly, there is a substantial change in the lead/lag relationship between 1961–1979 and 1980–1999. For the earlier period, there is a tendency for the SST pattern of mode 2 to appear 10 months after

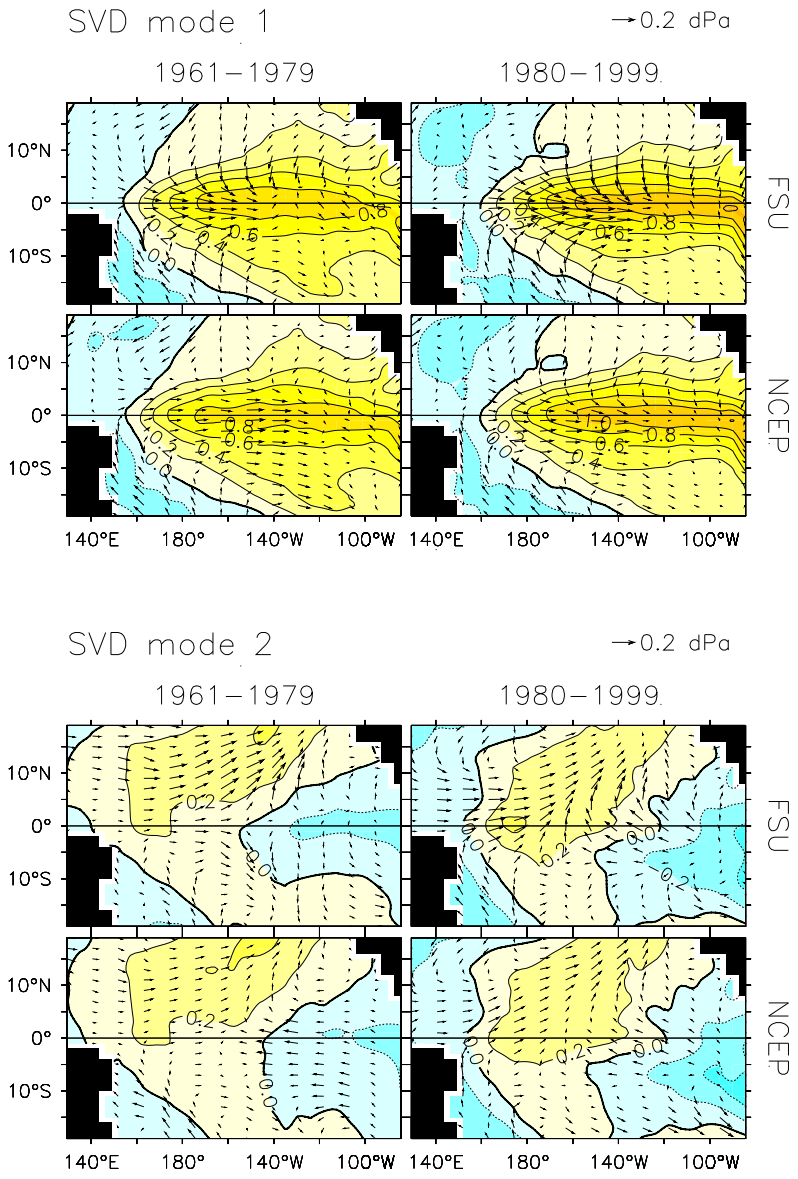

Figure 3.5: Leading singular vectors of the observed covariance of SST (°C) and wind stress anomalies (scale vector indicated).

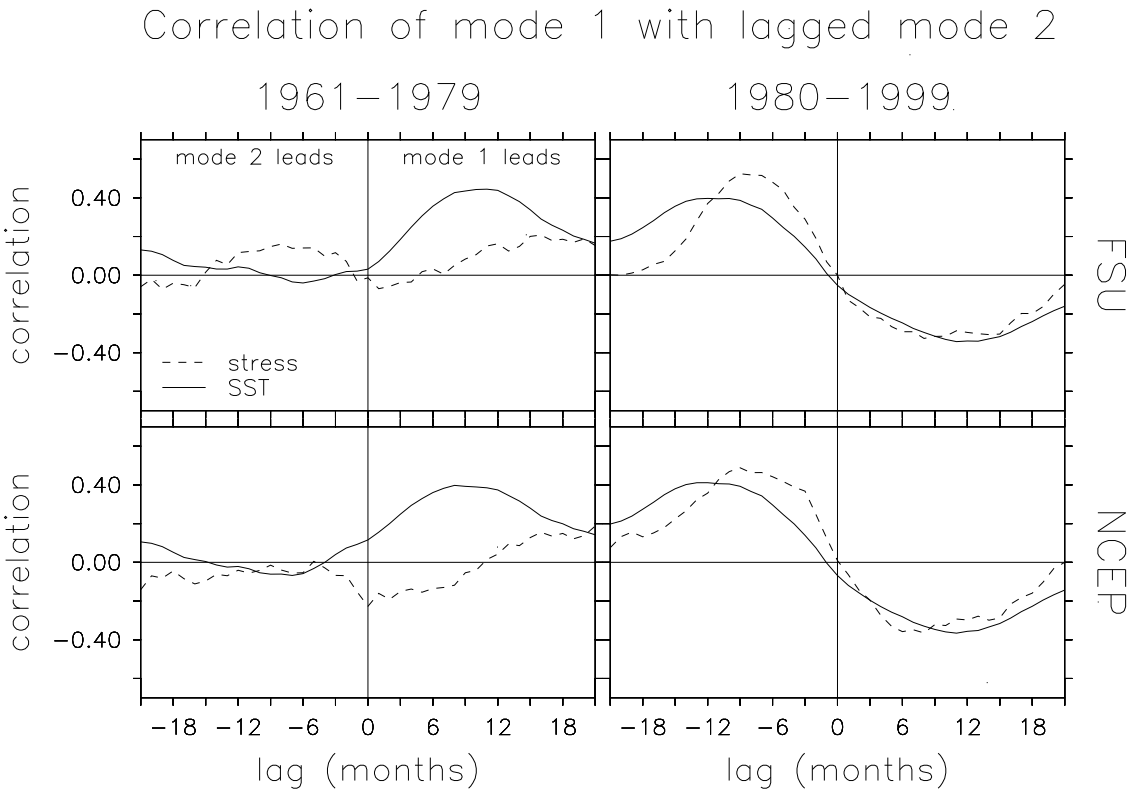

Figure 3.6: Correlation of singular vector expansion coefficients of mode 1 with those of lagged mode 2, for SST anomalies (solid) and wind stress anomalies (dashed).

mode 1, implying westward propagation of SST following the peaks of El Niño and La Niña. For the latter period, the propagation tendency reverses to eastward: mode 2 appears 11 months before mode 1, and again with opposite sign 11 months following mode 1, implying propagation of SST both before and after the peaks of ENSO events. This apparent change in the zonal propagation of SST anomalies in the late 1970s has been noted by Wang (1995) and Trenberth and Stepaniak (2001). It is not yet clear whether this signals a fundamental change in ENSO behavior; perhaps due to the brevity of the observational record, the change in propagation directions is not statistically significant (Harrison and Larkin, 1998). In any case, comparable changes in propagation direction appear to have occurred earlier in the 20th century (Trenberth and Stepaniak, 2001).

What is striking about Fig. 3.6 is that the propagation of SST anomalies during 1961–1979 is not mirrored in the wind stress over that period. During 1980–1999, on the other hand, there is clear southeastward propagation in the wind stress from the northwestern tropical Pacific to the central Pacific just south of the equator: mode 2 appears 8 months prior to mode 1 and again, with opposite sign, 9 months after mode 1.

The basic features of these correlations are similar between the datasets. While this is partly because the same SST product was used in both cases, it is reassuring to see that the primary lag relationship between the modes does not depend on which stress analysis is used.

We now have an empirical basis for the covariation of SST and wind stress, which captures a large fraction of the ENSO signal including the propagation of large-scale anomalies. We shall use this basis in the next section to construct a linear model of the stress anomalies.

3.4 Model for monthly-mean stress anomalies

On monthly time scales, tropical wind stress variability resembles a stochastic boundary-value problem. A reasonable linear model for the stress anomalies is

$$\mathbf{Y} = \mathbf{XW} + \mathbf{E} \tag{3.24}$$

where $\mathbf{W}$ is a $p \times q$ matrix of weights multiplying the instantaneous SST field, and $\mathbf{E}$ is an $n \times q$ matrix of shocks. A priori, we assume these shocks e_{lj} are $NID(0, \sigma_j^2)$; that is, normally and independently distributed in time, with zero mean and a variance that is stationary in time. The shocks may, however, be spatially correlated. We shall reexamine the validity of these assumptions once the model has been fitted.

3.4.1 Deterministic component

Our first task is to estimate the SST weights using the observations. Given our assumptions about the errors, we may use linear regression to help determine $\mathbf{W}$. There are not enough observations to regress each stress point directly onto all SST points, so we proceed by regressing the stress anomalies onto only the "coupled" part of the SST anomalies. Recalling that SVD provides an orthogonal basis that optimally captures the covariation of SST and wind stress anomalies, we truncate the mode expansion to filter

out SSTAs irrelevant to the covariation. A model for the stress data based on the leading N singular vectors of SST anomalies is

$$\mathbf{Y} = \widetilde{\mathbf{X}}_N^* \mathbf{R}_N + \mathbf{E}_N \tag{3.25}$$

where $\widetilde{\mathbf{X}}_N^*$ is the truncated $n \times N$ matrix of nondimensional SST singular vector expansion coefficients, $\mathbf{R}_N$ is the $N \times q$ matrix of weights, and $\mathbf{E}_N$ is the sum of the shocks and the truncation error.

Assuming the truncation error is $NID(0, \sigma_j^2)$, we can estimate the deterministic stress at each gridpoint by linear regression:

$$\widehat{\mathbf{Y}}_N \equiv \widetilde{\mathbf{X}}_N^* \widehat{\mathbf{R}}_N \tag{3.26}$$

$$\widehat{\mathbf{R}}_N \equiv \left(\widetilde{\mathbf{X}}_N^{*\prime} \widetilde{\mathbf{X}}_N^* \right)^{-1} \widetilde{\mathbf{X}}_N^{*\prime} \mathbf{Y} \tag{3.27}$$

where a "hat" tags matrices estimated from the data. An estimate of the shocks is given by the matrix of residuals

$$\widehat{\mathbf{E}}_N \equiv \mathbf{Y} - \widehat{\mathbf{Y}}_N \tag{3.28}$$

If our assumptions about the shocks and truncation error are valid, then the stress estimate (3.26) is *optimal* for the N modes in that it minimizes the sum of square residuals $\left\| \widehat{\mathbf{E}}_N \right\|^2$ over the data record.

Next we must select the number N of retained modes. To borrow from Einstein, things should be made as simple as possible, no simpler. We therefore build up the model in steps, adding predictors only if they clearly improve the simulation. To ensure that the stress has a smooth spatial dependence on SST, we regress the stress onto the same set of SST modes at all gridpoints. At each point, we quantify the misfit of the N-mode model using the sum of square errors (SSE):

$$\mathrm{SSE}_N = \sum_{l=1}^{n} (y_{jl} - \widehat{y}_{jl})^2 \tag{3.29}$$

where y_{jl} is the observation and $\widehat{y}_{jl}$ is the estimate, using the only the gravest N modes, of the wind stress at a particular gridpoint j and time l.

Now consider two models, one a regression onto N modes and the other a regression onto $N + k$ modes. For the regression model, it will always be the case that $\mathrm{SSE}_{N+k} < \mathrm{SSE}_N$. To check if this reduction in SSE is significant, we examine the ratio

$$f = \frac{(\mathrm{SSE}_N - \mathrm{SSE}_{N+k})/k}{\mathrm{SSE}_{N+k}/(n-1-N-k)} \tag{3.30}$$

Under the hypothesis H_0 that the extra k modes are useless as predictors, the numerator and denominator of (3.30) will be independent estimates of the true variance of the random shocks. In this case f will refer to an F distribution with k and $n-1-N-k$ degrees of freedom. We define the *significance* as the probability, given H_0, that f could lie as far as it does from unity. If the significance is small, then we may confidently reject H_0 and

Table 3.1: Number of modes retained in each statistical atmosphere model, with the percent of total square covariance, SST anomaly variance, and stress anomaly variance that project onto the singular vectors, and the percent of stress anomaly variance captured by the regression model.

	1961–1979		1980–1999	
	FSU	NCEP	FSU	NCEP
number of modes	4	5	3	5
% square covariance	97.7	98.7	97.6	98.1
% SSTA variance	81.2	85.2	76.3	80.7
% $\boldsymbol{\tau}'$ variance				
singular vectors	23.3	44.1	25.1	39.4
regression	12.6	23.0	13.9	19.9

conclude that the extra k modes (taken together) are useful predictors at the gridpoint of interest. We shall consider the change in SSE at a gridpoint to be significant when a two-tailed F test rejects H_0 with 99% confidence.

Using this definition for "significant model improvement" by a set of modes at a single gridpoint, we shall accept a set of modes only if they *significantly improve the model at more than half of the gridpoints*. One can then compute the appropriate number of modes to include using the following algorithm:

1) Begin with $N = N'' = 0$.

2) Determine the minimum number N' of additional SST modes required to produce a significant change in SSE at more than half of the stress gridpoints.

3) If $N' \leq 3$ then set $N = N + 1$, $N'' = N'$, and return to step 2.

4) Set $N = N + N''$ and stop.

This recipe ensures that predictors are included in the model if and only if they significantly improve the model fit, either by themselves or in a small group. Note that if the residual is not $NID(0, \sigma_j^2)$ then the F test will not be strictly valid, and the above procedure may add too many modes to the model. If the model assumptions are not severely violated, the problem can be addressed by increasing the gridpoint confidence level or increasing the number of gridpoints at which significant improvement is required. Otherwise a more sophisticated model may be necessary.

Following the recipe outlined above gives the models summarized in Table 3.1. These models account for a large fraction of the SSTA/$\boldsymbol{\tau}'$ covariance and the SSTA variance in the data, but account for less than half of the $\boldsymbol{\tau}'$ variance. Clearly most of the stress anomaly is not linearly related to SSTA; this part will be dealt with later.

The first two modes of the deterministic model are shown in Figs. 3.7–3.8 for each dataset and period (SST patterns are repeated from Fig. 3.5). These figures, plus those for higher modes, succinctly represent the entire linear model: projection of an SSTA onto these patterns produces the $\boldsymbol{\tau}'$ shown.

For mode 1, the wind stress patterns look very similar to regressions onto NINO3 SST anomalies (Chapter 2), since the SST predictor for this mode has a large loading in

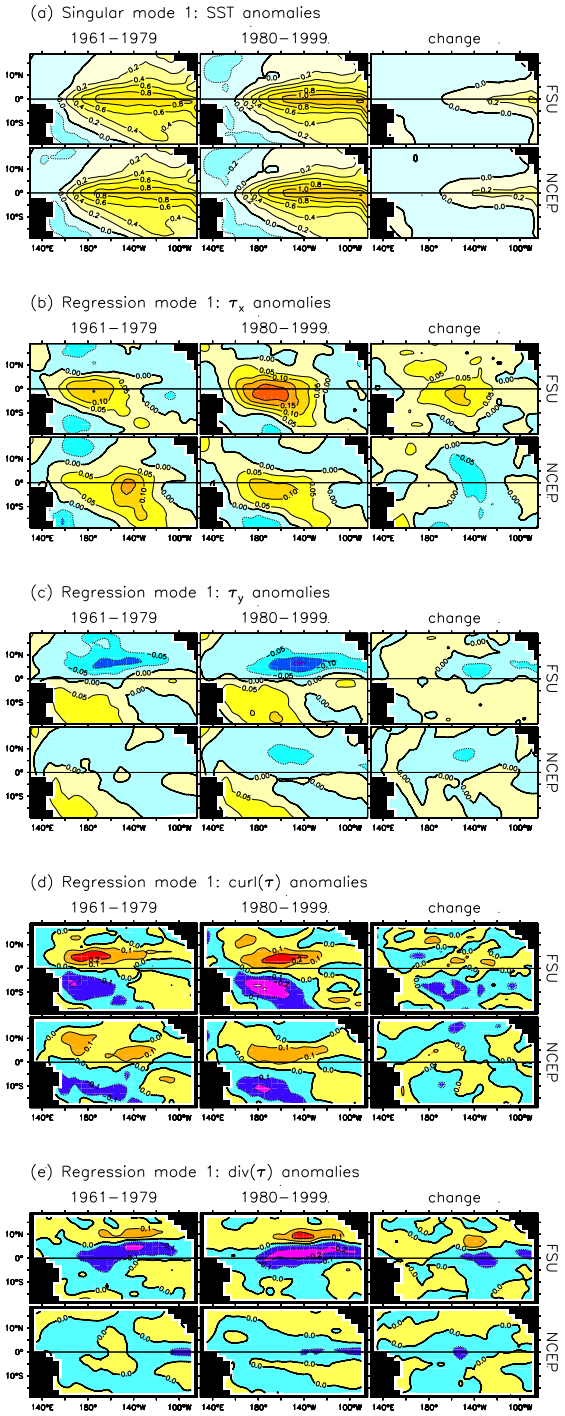

Figure 3.7: Mode 1 of the deterministic wind stress model. In this model, projection onto the (a) SSTA singular vector (°C) produces the corresponding anomalous (b) zonal stress (dPa), (c) meridional stress (dPa), (d) stress curl (0.001 dPa km^{-1}), and (e) stress divergence (0.001 dPa km^{-1}).

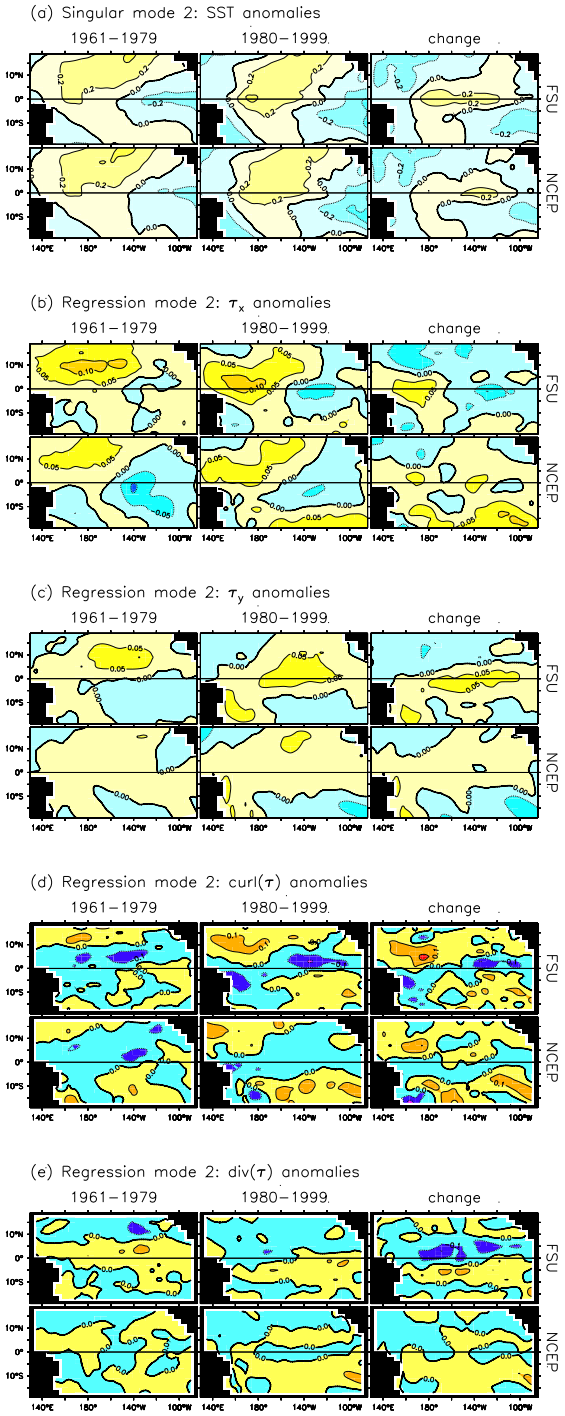

Figure 3.8: As in Fig. 3.7, but for mode 2 of the deterministic wind stress model.

the NINO3 region. The stress response consists of equatorial westerlies and off-equatorial easterlies in the west/central basin, weak easterlies near the eastern boundary, northerlies spanning the northern tropical Pacific, and southerlies across the southwest and equatorial southeast Pacific. There is cyclonic stress curl off-equator in the central basin, with the curl closer to the equator in the north than in the south. The stress response shows anomalous convergence equatorward of the ITCZ and SPCZ, and anomalous divergence poleward of these convergence zones.

There are several differences between the FSU and NCEP responses for mode 1. At the equator, the zonal stress response is weaker in the west Pacific in NCEP than in FSU; NCEP also shows a prominent peak near 145°W during 1961–1979 which is not evident in FSU. NCEP gives weaker meridional stress and convergence anomalies, especially in the vicinity of the climatological convergence zones. The two products also disagree as the nature of changes in the stress response between 1961–1979 and 1980–1999. FSU indicates a strengthening of the τ_x' response, with increased convergence in the equatorial eastern Pacific, while NCEP shows more of a weakening and a westward shift, with little change in the convergence in the east but instead increased convergence in the western and central Pacific.

The stress response for mode 2 is largely in quadrature with that for mode 1, partly due to the orthogonality constraint on the SSTA predictor pattern. For mode 2 there are westerlies in the northwest, and easterlies on and south of the equator in the east. Southerlies appear at and north of the equator, and northerlies appear in the southeast. The FSU and NCEP products again show large differences. The stress response for mode 2 is generally stronger in FSU, except during 1961–1979 when NCEP shows strong easterlies near 140°W. Between 1961–1979 and 1980–1999, FSU indicates a strengthening of the westerly response in the western equatorial Pacific and a weakening in the east, with stronger southerlies along the equator. NCEP shows somewhat smaller and less coherent changes.

By including mode 2 and higher modes, the statistical model can represent the gross wind stress response to SST anomalies, including large-scale propagating features. In summary, the deterministic part of the wind stress anomaly is obtained as follows:

1) Project the SST anomaly field $\mathbf{x}$ onto the leading N singular vectors of SST.

2) Divide through by the variance of the SST expansion coefficients to yield N predictors $\widetilde{\mathbf{x}}^*$.

3) Estimate the stress field from the predictors using the regression model.

In matrix form, the deterministic stress estimate is

$$\widehat{\mathbf{y}}_N = \mathbf{x}\mathbf{A}_N\mathbf{S}_{x^*}^{-2}\widehat{\mathbf{R}}_N \tag{3.31}$$

where $\mathbf{x}$, $\mathbf{A}_N$, and $\mathbf{S}_{x^*}$ have units of °C, and $\widehat{\mathbf{R}}_N$ has units of dPa (dynes cm^{-2}). A comparison of this model with a simpler version used by other authors is given in Appendix A.

3.4.2 Analysis of the residual

Having fit the deterministic model, we now ask to what extent the underlying assumptions are truly valid. In particular, we check whether the residual is really Gaussian noise with stationary variance, as assumed in the regression model. To do this, we obtain the residual stress by subtracting the deterministic model simulation from the data. This residual consists mostly of high-frequency components. It is quite possible that this "noise" plays an important role in the tropical climate system, since it accounts for a large fraction of the total stress variance (Table A.1).

Independence

The regression model assumes that the errors in the wind stress at each point are independent. To check this assumption we examine whether the residual is really uncorrelated in time. The autocorrelation of the residual stress at a lag of one month (not shown) is found to be rather large (exceeding 0.3) near the equator during 1980–1999.

Thus the residual at each point is not simply random from one month to the next. This suggests that the wind stress anomalies may not be a simple linear function of the large-scale SST anomalies. Possible sources of the autocorrelation in the residual include omitted explanatory variables (such as atmospheric moisture, which follows convection rather than SST), an incorrect functional form (e.g. a nonlinear rather than linear dependence on SST), or a combination of both (e.g. nonlinearity interacting with seasonal movements of the tropical convergence zones).

The lack of independence in the residual suggests that more information could be extracted by using a nonlinear or seasonally-dependent statistical model, or by including more predictor variables besides SST. It also indicates that the variance estimates used in the F-test for "significant model improvement" in the last section may be unreliable, leading to too many modes being accepted into the model. Fortunately, the estimates of the regression coefficients are unbiased even in the presence of serial correlation. Thus the model can still be useful, provided the number of accepted modes is sufficiently limited by the other criteria in the acceptance algorithm. Furthermore, our main goal is to correctly *partition* the deterministic and stochastic components and create realistic models for each. The mode selection algorithm must walk a tightrope between accepting too many modes, which would overfit the deterministic model, and accepting too few modes, which would waste information and possibly contaminate the stochastic model.

Normality

The second assumption of the regression model is that the residual at each gridpoint is normally distributed. To check this we use the powerful Shapiro and Wilks's W-test, Applied Statistics Algorithm AS 181 (Royston, 1982). The *significance* of this test is the probability that the test statistic, W, would be as extreme as observed if the data really were normally distributed. A small value of the significance at a gridpoint implies that the residual is probably non-normal at that point.

Significant non-normality is evident in both stress components (not shown), near the dateline and in scattered patches around the basin. This again suggests that a nonlinear

model of the stress may be justified. Two key nonlinearities could be included. The first is the quadratic dependence of the wind stress on wind speed, which in regions of strong trade winds tends to skew stress events toward extreme easterly values. The second is the zonal shift in anomalous trade winds observed during ENSO events, which is likely the result of the dependence of convection on thresholds of total SST. Zonal shifts in convection, which locally appear as nonlinearities, tend to skew the stress towards easterly stress extremes in the western Pacific, and westerly stress extremes in the central/eastern Pacific.

Stationarity

To test whether the variance is constant through time at each gridpoint, we must check whether subsets of the data have variances which are consistent with the overall variance of the dataset. In particular we examine the residual zonal stress averaged over the equatorial band, 2°S–2°N.

To test whether the residual variance is a function of the annual cycle, we bin the data by calendar month, and compute the ratio of monthly binned variances to the overall variance (Fig. 3.9). In each case, the local significance test rejects substantially more than 1% of the points at the 0.01 level. If the null hypothesis were true, the chance of this many local tests being rejected would be small: Fig. B.1 indicates this chance would be less than 5% (except for FSU 1980–1999, which would be less than 30%). This suggests the residual variance is not constant throughout the year at many gridpoints. In the western central Pacific, the variance is particularly strong during boreal winter and spring, and weak during boreal summer and fall.

To check whether the variance changes slowly through time, we bin the data by year and calculate the variance of the data in each of the bins. The ratio of yearly binned variances to the overall variance is shown in Fig. 3.10. Contours surround points where the hypothesis of stationarity of the residual variance is rejected at the 0.01 level, according to a two-tailed F test. Again, the local significance test rejects substantially more than 1% of the points at the 0.01 level. The probability of this many local tests being rejected, if the null hypothesis were true, would be small (less than 5%, Fig. B.1). This suggests that the residual subannual variance is nonstationary from year to year at some of the gridpoints.

Nonstationarity will cause the regression to overfit the noisy periods during boreal spring and the 1960s and 1980s, and underfit the quiet periods during boreal summer and the late 1970s and early 1990s. It is possible that the nonstationarity and non-normality of the residual are related, through some nonlinearity of the wind stress dependence on SST. If this is the case, then a nonlinear transformation applied to the data might improve the model. In the interest of simplicity, however, we shall proceed with the fully linear version, having noted its problems. While significant, the nonstationarity of the residual variance is probably not strong enough to greatly affect the regression model. The model fit should still be accurate, since the SST predictors are only slightly skewed.

Lag correlation with SST anomalies

As noted by Syu and Neelin (1995), the deterministic model acts like a low-pass filter for the wind stress, passing only variability linearly coupled to the slowly-varying SST field.

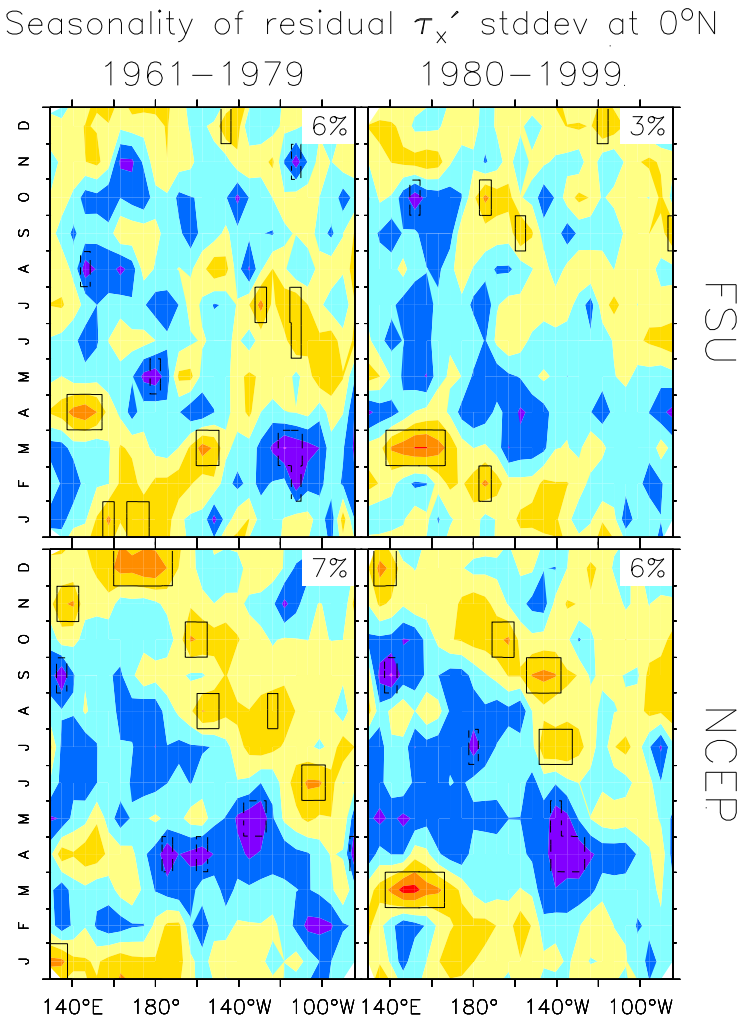

Figure 3.9: Ratio of binned monthly standard deviation to the standard deviation of the entire residual τ_x' timeseries at the equator. Warm colors indicate above-average variance, cool colors below-average variance. Contours surround points where the hypothesis of stationarity of the residual variance is rejected with 99% confidence, according to a two-tailed F test. Points within solid and dashed contours are rejected for having too much and too little variance, respectively. Labels at the upper-right indicate the fraction of points rejected by the F test.

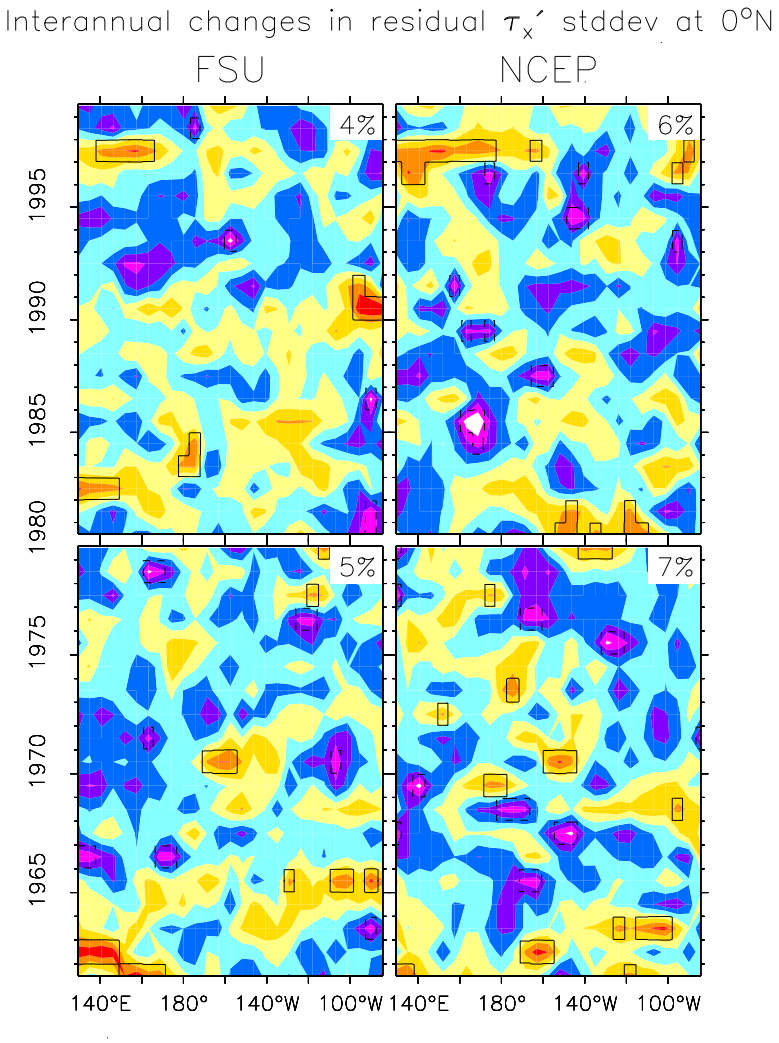

Figure 3.10: Ratio of running annual standard deviation to the standard deviation of the entire residual τ_x' timeseries at the equator. Otherwise as in Fig. 3.9.

Some low-frequency components are present in the residual, however. In particular, the model tends to underestimate the westerly stress response to the strong El Niño event of 1982–83, and misses the easterly stress events in the western Pacific during 1966-67 and 1983–84. It is possible that nonlinearities in convective heating and wind stress emerge at these times, leading to coupling between the residual stress and SST.

In the coupled context, the noisy component of the stress can feed back on the deterministic stress by affecting SST. An unusually strong random westerly wind burst in the west Pacific, for example, would tend to induce an eastward advection of the warm pool and a deepening of the equatorial thermocline in the eastern basin. As a result, after some lag the equatorial SST would tend to increase, the wind stress would respond "deterministically" with a westerly wind anomaly over the central Pacific, and the coupled system could amplify the perturbation into a warm event. One would therefore expect residual stress anomalies to lead SST anomalies, by some characteristic time associated with the transient growth in the coupled system.

Fig. 3.11 shows how equatorial SST anomalies are linked to the residual zonal stress in the western Pacific. The lag correlations are rather weak, since the residual is noisy, but they are statistically significant (at the 0.05 level) over large regions of the diagram. Although Fig. 3.11 indicates that residual wind bursts in the west tend to accompany reversals of the SST-induced stress over the entire basin, it says nothing about cause and effect. One plausible hypothesis is that stochastic wind events in the west Pacific help to bring about ENSO transitions. An alternative hypothesis is that during the transitions between events, the stress anomalies in the west Pacific are more nonlinear, or depend more on smaller-scale SST anomalies. Without dynamical modeling experiments or further physical evidence, it would be difficult to decide which hypothesis is correct.

Nevertheless it is clear that the model and residual stresses are linked in time. In all four cases, westerly stress events in the western equatorial Pacific are preceded 6 to 18 months earlier by cold SSTA in the central Pacific, and followed 2 to 12 months later by warm SSTA in the central and eastern Pacific. The linkage is stronger for FSU than NCEP, and is stronger for 1980–1999 than for 1961–1979. There is tantalizing evidence that the residual stress in the west is related to small-scale anomalies in the zonal SST gradient at zero lag: all cases show anomalously warm SST east of the reference point, and anomalously cold SST west of the reference point. One interpretation is that local coupled feedbacks in the west Pacific amplify random small-scale stress anomalies; these amplified anomalies could then induce Kelvin wave disturbances which give rise to ENSO events through large-scale feedbacks. This picture is consistent with emerging views of possible links between the Madden-Julian Oscillation and ENSO (Zhang et al., 2001a).

3.4.3 Stochastic model

It is important to note that "not knowing the properties of the noise" is very different from "not having any noise." Various types of noise can sustain a damped ENSO, giving rise to sustained interannual variability, whereas in the absence of noise such a system would be quiescent. Stochastic forcing not only feeds energy into a system but also acts as a sort of dynamical averager—stirring up variability that visits a broader range of system dynamics, with random excursions away from asymptotic fixed points and limit cycles. A

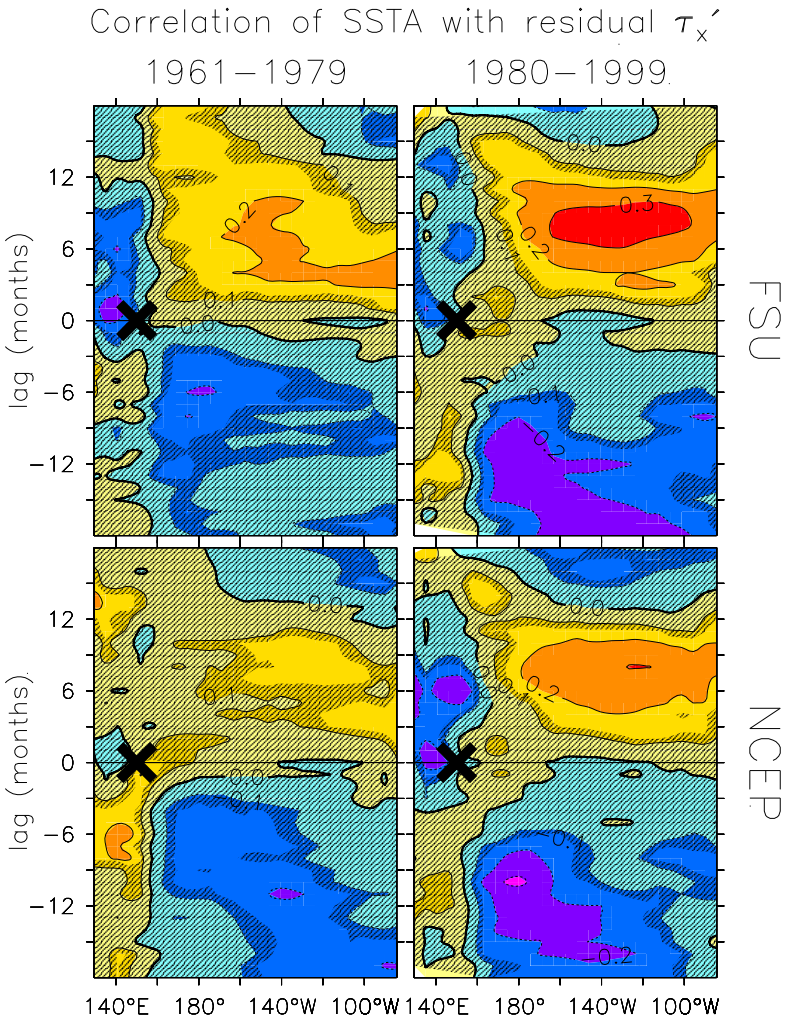

Figure 3.11: Correlation of the lagged equatorial SST anomaly with the residual zonal stress anomaly at a point (marked) on the equator at 150°E. Hatching masks values exceeded by more than 5% of correlations between white noise and the SST anomaly.

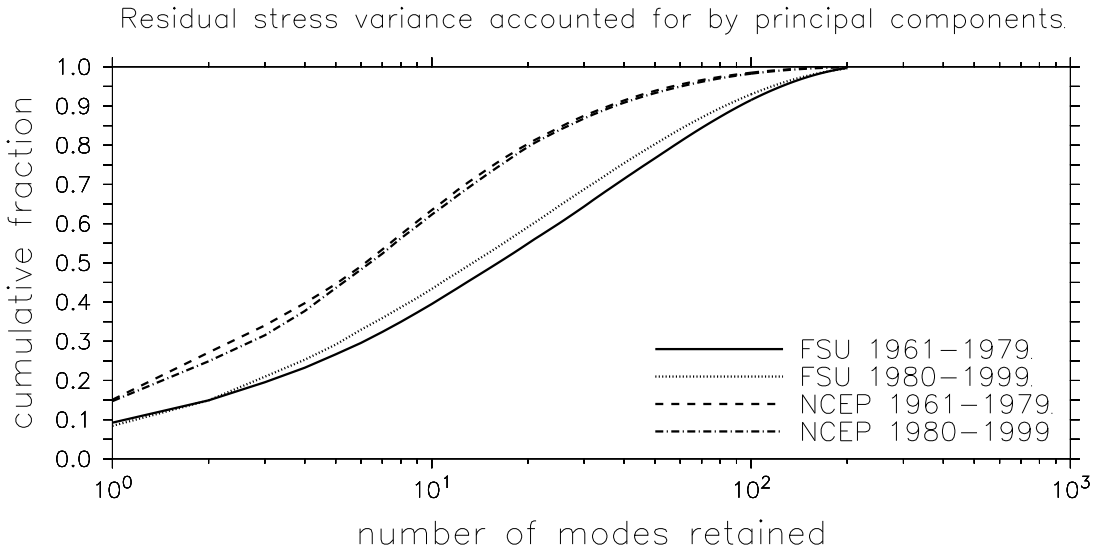

Figure 3.12: Cumulative fraction of residual variance of monthly-mean tropical Pacific SST and wind stress, accounted for by the principal components of the residual. The residual consists of the original data minus the deterministic model estimate.

model with stochastic forcing may also exhibit different sensitivity to parameter changes, since bifurcations and linear modes get smeared out and slow transients begin to dominate the system dynamics.

For these reasons, having a simple, gross estimate of the noise may be better than having no noise at all. Although part of the residual stress in the data may have arisen from observational and analysis errors, there is clearly subannual-scale variability present in the tropical climate system, and while the exact form of the noise is uncertain, one can still model gross aspects like its variance and spatial patterns. In lieu of a more sophisticated model of the residual, we shall approximate it as autocorrelated ("red") noise.

To account for spatial correlations in the residual stress, we decompose $\widehat{\mathbf{E}}_N$ into its principal components using (3.23):

$$\widehat{\mathbf{E}}_N = \widetilde{\widehat{\mathbf{E}}}^*_{N,\widehat{e}} \mathbf{B}'_{N,\widehat{e}} \tag{3.32}$$

The cumulative fraction of residual variance accounted for by the principal components is shown in Fig. 3.12. The fact that the residual has so many degrees of freedom is consistent with the hypothesis that it is mostly noise. The principal components account for the residual variance more efficiently for NCEP than for FSU. For NCEP, over half the variance is accounted for by 7 modes, 90% by 40, and over 99% by 120. For FSU, over half the variance is accounted for by 17 modes, 70% by 40, and over 99% by 180. For NCEP there is little difference in the representation efficiency between 1961–1979 and 1980–1999, while for FSU the efficiency is somewhat greater during the latter period.

The first three PCA patterns of the residual are shown in Fig. 3.13. Although these patterns account for a relatively small fraction of the total residual variance, they are

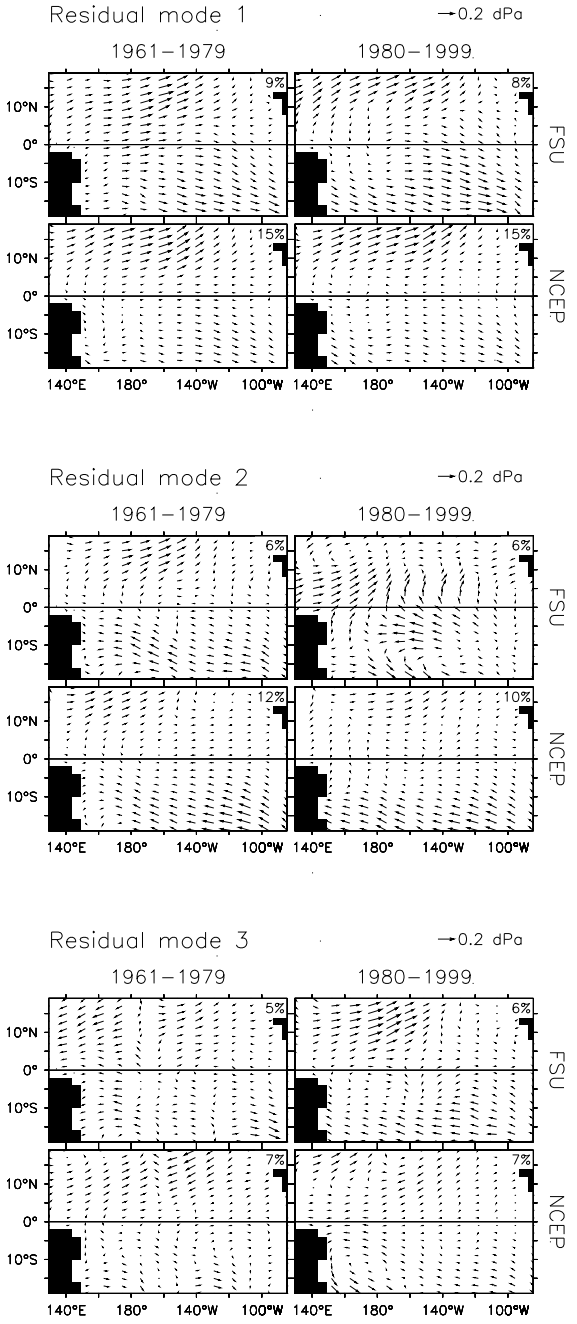

Figure 3.13: Leading principal components of the wind stress residual. The residual consists of the total stress anomaly minus the stress anomaly predicted by the deterministic model. Labels at the upper-right indicate the fraction of residual variance associated with each pattern.

dynamically important because they represent large-scale changes that could affect SST. Most of the patterns show only a weak equatorial signature, but they are associated with changes in the near-equatorial stress curl which could induce equatorial Rossby waves. These waves could then affect SST through equatorial zonal advection, and through reflection into Kelvin waves at the western boundary.

Residual mode 1 represents an overall change in the strength of the trade wind stress, especially north of the equator. The patterns differ between FSU and NCEP, with FSU showing a larger signature in the equatorial and southern tropical Pacific. The patterns also change between 1961–1979 and 1980–1999: FSU shows stronger loading in the off-equatorial western Pacific, with weaker loading in the equatorial western Pacific; similar changes occur in NCEP but to a lesser extent. Residual mode 2 represents more of an antisymmetric change about the equator in the off-equatorial zonal stress. The patterns look fairly similar except for FSU 1980–1999, which shows a much stronger signature near the equator and in the central Pacific. Subsequent modes are heavily constrained by orthogonality, but retain large-scale structures (e.g. residual mode 3).

Although the residual stress modes are coherent in space, their observed projection coefficients are fairly random. Examination of the autocorrelation functions for the residual PCA expansion timeseries indicates that beyond 1-month lag, the autocorrelation of the monthly-mean residual stress is below 0.2 for all modes. The cross-correlations of the PCA timeseries (not shown) are also small at all lags; indeed, the simultaneous cross-correlations are identically zero by (3.20). Thus to an good approximation, the PCA expansion coefficients are random from one month to the next.

We therefore assume that the residual PCA expansion timeseries evolve like red noise, with decorrelation times shorter than one month. Let $\mathbf{f}(t)$ represent the set of expansion coefficients at time t. If for simplicity we assume all the coefficients have the same decorrelation time, the expansion timeseries evolve according to

$$\mathbf{f}(t) = \mathbf{f}(t - \Delta t)e^{-\zeta} + \mathbf{g}(t)\sqrt{1 - e^{-2\zeta}} \tag{3.33}$$

where Δt is the time step and $\mathbf{g}(t)$ is a random vector whose elements are $NID(0, 1)$; that is, normally and independently distributed with zero mean and homogeneous unit variance. The ratio ζ of the time step to the decorrelation time determines the "color" of the noise. As ζ decreases, the noise spectrum becomes redder, and the noise timeseries take more extended departures from zero. For a decorrelation time of 1 week and a timestep of 1 month, $\zeta \approx 4$ and the noise is practically white in time, $\mathbf{f}(t) \approx \mathbf{g}(t)$, as required by the data.

In summary, the noise residual is modeled using

$$\widehat{\mathbf{E}}_N \approx \mathbf{F}\mathbf{B}'_{N,\widehat{e},M} \tag{3.34}$$

Each column of the $n \times M$ matrix $\mathbf{F}$ is a noise timeseries with a decorrelation time of one week. The $q \times M$ matrix $\mathbf{B}_{N,\widehat{e},M}$ consists of the leading M principal patterns of the residual stress estimated from the leading N singular modes.

While the PCA patterns efficiently represent the variance of the residual, it is possible that the ocean will respond to only part of the noise signal represented by these patterns.

One could imagine using a different basis for the noise, such as a set of "optimal perturbations" which produce maximal growth of SST anomalies over some specified time window (Houtekamer, 1995; Penland and Sardeshmukh, 1995; Moore and Kleeman, 1996, 1997a,b, 1998). Such optimal patterns would be directly relevant to the observed coupled system over recent past. On the other hand, they might not remain optimal in the future if the dynamics or climate of the ocean changes. Furthermore, optimals can differ between reality and models, as well as among different models (Moore and Kleeman, 2001). For generality, therefore, we choose here to model *all* the noise, represented efficiently by the PCA basis, and let subsequent coupled systems select the optimal patterns for themselves.

3.4.4 Full model

Combining (3.31) and (3.34), a model for the full wind stress anomaly field at time t is

$$\mathbf{y}(t) = \mathbf{x}(t)\mathbf{A}_N\mathbf{S}_{x^*}^{-2}\widehat{\mathbf{R}}_N + \mathbf{f}(t)\mathbf{B}'_{N,\widehat{e},M} \tag{3.35}$$

which simply says that stress anomalies depend on certain preferred patterns of SST anomalies and noise. The deterministic part of the stress is generated by a regression $\widehat{\mathbf{R}}_N$, estimated from historical data, which relates stress anomalies to dominant coupled SST/stress signals. These nondimensional signals arise from projecting the current SST anomaly $\mathbf{x}$ onto the leading N singular vectors $\mathbf{A}_N$ of historical SST/stress anomalies, and then dividing by the historical SST variance $\mathbf{S}_{x^*}^2$. The stochastic part of the stress is generated using random combinations $\mathbf{f}(t)$ of the principal "noise" patterns $\mathbf{B}'_{N,\widehat{e},M}$, which are estimated from the part of the historical data not explained by the deterministic model.

3.5 Discussion

The statistical model (3.35) reproduces many of the observed characteristics of tropical Pacific wind stresses. Analysis of the model errors indicates that the simulation is good, but may not be a complete description of the stress. Thus a few caveats are in order.

The deterministic part of the model depends on large-scale SST anomalies only. In reality, tropical surface winds are forced by pressure gradients which depend on gradients in atmospheric heating. Part of the heating is related to local SST, since conduction, mixing, and radiation convey surface temperatures upward into the atmospheric boundary layer. This local heating then helps determine the locations of large-scale convergence, where further heating is achieved through condensation. Since water vapor is supplied by evaporation over a wide area, part of the local heating depends on nonlocal effects. The present model will capture many of these nonlocal effects, but will miss those that arise from temperature variability outside of the tropical Pacific domain. The deterministic model also neglects effects of local, small-scale SST anomalies on the local circulation; these effects are parameterized by the stochastic model.

The deterministic model is linear. Surface wind stresses, however, depend nonlinearly on the wind; often quadratic dependence is assumed. The winds, in turn, depend nonlinearly on the atmospheric heating, since advection and mixing can play roles in the

momentum balance. Positive feedbacks between atmospheric heating and moisture convergence can also produce nonlinearity, since regions of warm SST and mean convergence have a heightened sensitivity to SST anomalies; such a feedback has been incorporated into simple dynamical atmospheric models like that of Zebiak (1986). The present statistical model can only hope to represent the linear part of the above processes. Fortunately, since there is a large amount of data at each gridpoint and the skewness of the residual is not very severe, we expect that the linear regression coefficients will be little affected by the non-normality. While the deterministic model may not completely describe the stress, we can be reasonably confident that it is at least a meaningful *linear* fit to the data.

The model is based on an assumption that the SST/stress relationship and the noise statistics are stationary in time. Analysis of the residual stress indicates that this assumption is not strictly correct: the noise variance changes through the calendar year, and also from year to year. To handle the seasonal nonstationarity, one could block the seasonal effects by fitting a separate model to each month of the calendar year. This time-dependent model would suffer from a substantial reduction in the regression degrees of freedom, but could also benefit from reduced noise variance for each month. It is difficult to statistically quantify the effects of interannual changes on the SST/stress relationships and noise variance, given the present data record. For now, we shall content ourselves with simulating only the most robust, linear part of the wind stress over the entire period 1961–1999.

By removing the seasonal cycle from the data before performing the model fit, we eliminate the assignable seasonal influence of solar forcing on the atmosphere. This is ideal for simulation of anomalies from the seasonal cycle, since it reduces the variance of the data and thereby improves the model fit. It would be easy to repeat the model-fitting exercise using departures from the annual-mean climatology instead of anomalies from the mean seasonal cycle (Syu and Neelin, 1995). Though such a model would be less skillful at predicting anomalies from the seasonal cycle, it would allow simulation of the seasonal cycle itself as well as anomalies from it.

The statistical model is "trained" to respond only to SST anomalies which have analogues in the observed record. It is possible that forcing the statistical model with very unusual SST anomalies could generate an unphysical stress response, which might preclude highly idealized experiments involving observationally rare SST patterns. On the other hand, experience with so-called "constructed analogues" (van den Dool, 1994; Landsea and Knaff, 2000) suggests that skill can be achieved so long as the predictor pattern is representable as some linear combination of observed patterns. Thus it is the stress model may be able to produce a reasonable response even when presented with unusual SST anomalies.

A statistical model is only as good as the data used to build it. If the data has systematic biases, these biases will also show up in the statistical model. Trends were not removed from the data and so have been incorporated into the model. This is desirable if the trends are real, since they provide additional information on the relationship between SST and wind stress to be used by the linear regression. However, if the trends are due only to changes in data density, measurement procedures, or quality control, the model will be a fit to this spurious variability and not to the real system. The primary difficulty is choosing the wind stress product with which to build the model; a consensus product is much needed

to bridge the large gaps between these analyses. Until such a product is available, it may be wise to examine the sensitivity of diagnostic results and model experiments to different wind products. If a statistical model is used, it should not necessarily be viewed as ground truth, given the large uncertainties in the data.

The key to future progress in this area may lie in decoding and successfully modeling the structure of the noise. We have shown that there is some information left in the residual, including spatial and temporal coherence that may be dynamically important for ENSO. More sophisticated stochastic models may be better able to extract this information, which would enhance understanding and modeling of intraseasonal variability and its interactions with ENSO.

Having noted the caveats above, the author believes the statistical model provides a reasonable baseline for simulation of tropical Pacific wind stress anomalies. The model developed here (specifically FSU 1980–1999) will be used in the next chapter to build a fully coupled model of the tropical Pacific climate system.