Regime-Dependent Autoregressive Time Series Modeling of the Southern Oscillation

FRANCIS ZWIERS

Canadian Climate Centre, Downsview, Ontario, Canada

HANS VON STORCH

Max Planck Institut für Meteorologie, Hamburg, Federal Republic of Germany (Manuscript received 14 November 1989, in final form 9 May 1990)

ABSTRACT

The class of "regime dependent autoregressive" time series models (RAMs) is introduced. These nonlinear models describe variations of the moments of nonstationary time series by allowing parameter values to change with the state of an ancillary controlling time series and possibly an index series. The index series is used to indicate deterministic seasonal and regimal changes with time. Fitting and diagnostic procedures are described in the paper.

RAMs are fitted to a 102-year seasonal mean tropical Pacific sea surface temperature index time series. The models are controlled by a seasonal index series and one of two ancillary time series: seasonal mean Adelaide sea level pressure and Indian monsoon rainfall, which have previously been identified as possible precursors of the extremes of the Southern Oscillation (SO).

Analysis of the fitted models gives clear evidence for the seasonal variation of the statistical characteristics of the SO. There is strong evidence that the annual cycle of the SO index depends upon the state of the SO as represented by the ancillary time series. There is weaker evidence which suggests that its autocorrelation structure is also state dependent.

1. Introduction

There are examples of processes that have season and state dependent statistical properties in many areas of geophysics and it is reasonable to suspect that the Southern Oscillation (SO) is one such phenomenon. Southern Oscillation "warm" and "cold" events are clearly linked to the seasonal cycle (Rasmusson and Carpenter 1982; Wright et al. 1988). The existence of numerous studies that describe possible SO precursors, including van Loon and Shea (1985), Wright et al. (1988) and Barnett et al. (1989), also suggests that the variability of the SO is state dependent. The purpose of this paper is to investigate a class of nonlinear time series models that can accommodate both seasonal and state dependence in both their mean and stochastic properties and to illustrate their use by applying the models to a well-known SO index, specifically Wright's (1984) SST index.

There have been numerous studies of the stochastic behavior of the SO. Wright (1977, 1985, 1988) examined the relationship between the Southern Oscil-

Corresponding author address: Dr. Francis W. Zwiers, Numerical Modeling Division, Canadian Climate Center, Atmospheric Environment Services, 4905 Dufferin Street, Downsview, Ontario, Canada M3H 5T4.

lation index and other meteorological and oceanographic parameters and made some efforts to quantify the seasonal march of the SO autocorrelations. Chu and Katz (1985, 1987, 1989) used stationary and seasonal ARMA models to describe the statistics of the Southern Oscillation index and to characterize the stochastic behavior of the SO with some success. Barnett (1983, 1984a,b, 1985) and Gutzler and Harrison (1987) used complex EOF analysis techniques. In this paper we describe the use of nonlinear time series, models.

A few papers have been published on the question of whether the SO is a chaotic system. Fraedrich (1988), in a study of the SO's predictability, found no clear support for the existence of a strange attractor. Hense (1987), however, using different data than Fraedrich, found evidence for the existence of an attractor of fractal dimension between 2.5 and 6.

Time-domain time-series analysis tools such as the ARIMA (autoregressive integrated moving average) models of Box and Jenkins (1976), used by Chu and Katz (1985, 1987, 1989) and others to analyze SO index time series, are able to describe seasonal variations in stochastic structure and certain kinds of nonstationary stochastic behavior, but are not able to represent processes in which the statistical properties depend on the seasonal cycle and on the state of related processes

or on the basic state that is exposed to stochastic variations.

However, when screening the recent statistical literature we found a relatively large class of nonlinear time series models ("Threshold Autoregressive Models": Tong 1983), which we thought might be useful for the analysis of geophysical processes affected by a stochastic environment. This class of models, which we call "regime dependent autoregressive models" (RAMs), forms an extension of the usual class of linear autoregressive (AR) time series models (Box and Jenkins 1976) that have been applied extensively in the geophysical literature. The difference between RAMs and ordinary AR models is that in a RAM, two external variables are employed to control the choice of autoregressive parameters used at a particular time. One of the external variables is an indicator describing the seasonal march. The second external variable is usually another stochastic time series. When the latter crosses a threshold, the parameters governing the evolution of the modeled time series change. Tong shows that the class of models is quite general in nature and that it is able to accommodate a rather large range of nonlinear characteristics.

Fitting RAMs is computationally straight forward because of their "piece-wise" linear nature. Moreover, it is possible to objectively choose the "best" fitting RAM from a group of such models by means of an information-theory-based criterion such as the Akaike information criterion (AIC) [Akaike 1973]. Despite these positive characteristics, which are not shared by other classes of nonlinear time series models, we have learned from our modeling experience that these models must be fitted with some care and must be interpreted carefully.

In this study we identify RAMs that describe the nonstationary behavior of the Southern Oscillation (SO) as it is represented by Wright's homogenized SST index (Wright 1984). We will refer to this index as the SO_{SST} index to distinguish it from the widely used SOI index which is based on the normalized pressure difference between Darwin and Tahiti. The fitted models use two indices that have previously been identified as possible precursors of the extremes of the SO as controlling stochastic time series; the zonal pressure difference across the southwest Pacific (van Loon and Shea 1985, 1987; Xu and Storch 1990) and the Indian monsoon rainfall (Barnett et al. 1989; Elliot and Angell 1987).

A description of the RAM class of models is given in section 2. This is followed by a description of the fitting procedure in section 3. An application of RAMs to the SO_{SST} index using both southwest Pacific sea level pressure (SLP) and Indian monsoon rainfall indices as ancillary time series is described in section 4. Results of this analysis are described in section 5, and the paper is concluded with a discussion and summary in section 6.

2. Regime dependent autoregressive time series models (RAMs)

a. Formal definition of a RAM

A Regime dependent Autoregressive Model (RAM) X_t controlled by an indicator series S_t and an ancillary time series Y_t is described by a stochastic difference equation of the form

$$X_{t} = a_{0}^{s(t),y(t)} + a_{1}^{s(t),y(t)}X_{t-1} + a_{2}^{s(t),y(t)}X_{t-2} + \cdots + a_{n}^{s(t),y(t)}X_{t-n} + \epsilon_{t}.$$
 (1)

We refer to p as the "order" of the RAM. Superscripts s(t) take integer values $1 \le s(t) \le s$ to represent the seasons indicated by S_t . The latter is deterministic and strictly periodic with period s. Superscripts y(t) identify the "region" occupied by the ancillary time series Y_t at a lag of d seasons. These regions are defined in terms of "threshold parameters" $-\infty < T_1 < \cdots < T_p < +\infty$ and the "delay parameter" d by

$$y(t) = j$$
 if $Y_{t-d} \in (T_{j-1}, T_j],$
 $j \in \{1, \dots, \varphi + 1\}$ (2)

where $T_o = -\infty$ and $T_{y+1} = +\infty$. We refer to the combination of a particular season s and region j as a "regime" because the RAM has constant stochastic properties for all observing times t that fall into a given season and for which the ancillary time series takes values within a given region. The indices s(t) and y(t) define a total of s(y+1) possible different regimes for the modeled time series X_t . The "innovation process" ϵ_t is a zero mean Gaussian white noise process whose variance in regime (s,j) is $\sigma_{s,j}^2$. The RAM coefficients $a_i^{s,j}$ depend upon the regime (s,j).

b. Model hierarchy

The general definition given by (1) and (2) incorporates a hierarchy of models of increasing complexity. At the bottom of this hierarchy there are traditional AR models that ignore the seasonal variation and external forcing. Disregarding the ancillary time series. or setting $\psi = 0$, leads to the second level in the hierarchy, the ordinary seasonal AR models. The third level in the hierarchy, first-moment RAMs, is obtained by allowing only the constant term coefficient a_0 of the seasonal AR model to depend on the ancillary time series. In first-moment RAMs, the ancillary time series does not affect second moment statistics such as the variance. The top of the hierarchy, full RAMs, is obtained by using the full version of (1) and (2) that describes a seasonal model in which both the annual cycle and variation about the annual cycle is modulated by the ancillary time series.

In our example, presented in sections 4 and 5, we restrict ourselves to the upper three levels of the hierarchy and we will refer to these three levels as the RAM Hierarchy.

c. j-Moments

Clearly, a RAM is in general a nonstationary model and, thus, traditional moments can not be defined. It is, however, possible to derive moments conditional upon a given scenario of regimes. In particular, we define "conditional j-moments" to be the s(t)-dependent moments of a RAM that remains in region $j \in \{1, \dots, y + 1\}$ for an infinitely long period of time.

1) CONDITIONAL j-MEANS

The "conditional j-means," $E(X_t|y(t)=j; j=t, t-1, \cdots)$, of a RAM are given by the solution of a system of linear equations to be described below. We will simplify the notation by writing $E(\cdot|j)$ to indicate $E(\cdot|y(t)=j, j=t, t-1, \cdots)$. By taking expectations on both sides of (1) we find

$$E(X_t|j) = a_0^{s(t),j} + a_1^{s(t),j} E(X_{t-1}|j) + a_2^{s(t),j} E(X_{t-2}|j) + \cdots + a_n^{s(t),j} E(X_{t-n}|j).$$
(3)

This system of equations may be closed by noting that $E(X_t|j)$ is periodic with period s when y(t) = j for all t.

CONDITIONAL j-VARIANCES AND j-COVARIANCES

The "conditional *j*-second moments" may be obtained by writing the model in its infinite moving average (MA) form:

$$X_{t} = E(X_{t}|j) + \epsilon_{t} + c_{1}^{s(t),j} \epsilon_{t-1} + c_{2}^{s(t),j} \epsilon_{t-2} + \cdots$$
(4)

which we assume to be mean square convergent. Using the independence properties of the innovations we see that

$$\sigma^{2}(X_{t}|j) = \sigma_{s(t),j}^{2} + (c_{1}^{s(t),j})^{2} \sigma_{s(t-1),j}^{2} + (c_{2}^{s(t),j})^{2} \sigma_{s(t-2),j}^{2} + \cdots$$
 (5)

The lag Δ "j-autocovariance" is also easily derived. We have that

$$Cov(X_t, X_{t+\Delta} | j) = E[(\epsilon_t + \sum_{u=1}^{\infty} c_u^{s(t), j} \epsilon_{t-u})$$

$$\times (\epsilon_{t+\Delta} + \sum_{v=1}^{\infty} c_v^{s(t+\Delta), j} \epsilon_{t+\Delta-v})]$$

and, thus,

$$Cov(X_t, X_{t+\Delta}|j)$$

$$= c_{\Delta}^{s(t+\Delta),j} \sigma_{s(t),j}^2 + \sum_{i=1}^{\infty} c_1^{s(t),j} c_{i+\Delta}^{s(t+\Delta),j} \sigma_{s(t-i),j}^2.$$
 (6)

Note that the *j*-autocovariance function is not a simple function of the time difference alone. It depends

upon both lag Δ and time t. In fact, there is a distinct autocovariance function for each season which describes the linear relationships between the present realization of X_t and realizations of the time series at other times.

To calculate the *j*-second moments of a RAM it is necessary to derive the MA coefficients $c_i^{s,j}$ from the AR coefficients $a_i^{s,j}$. To do so, the past states X_{t-1} , X_{t-2}, \ldots , on the right-hand side of Eq. (1) are replaced by linear combinations of previous states by recursive application of (1). This recursion results in a sequence of gradually improving finite MA approximations of (4). The necessary calculations are easy to implement on a computer. It is, of course, impossible to exactly represent the infinite moving average by computational means because the recursion must be stopped after a finite number of steps. Thus, (4) must be approximated by a finite MA process of some order. A suitable stopping criterion can be based on the variance of the finite MA approximation obtained after every iteration; the approximation will be satisfactory when its relative change becomes small.

3. Model fitting procedure

One of the great advantages which RAMs have over other nonlinear time series models is that they are easy to fit to data using approximate maximum likelihood estimators (MLEs). Exact MLEs of RAM parameters can be obtained but require the use of nonlinear minimization techniques. Approximate MLEs may be easily obtained via least squares and the resulting parameter estimates enjoy the same asymptotic properties as exact MLEs. We therefore decided to use the approximate procedure described by Tong (1983).

a. Model fitting and identification

Model fitting is done in two steps. First, it is assumed that the delay parameter d and the threshold set $\mathbb{T} = \{T_1, T_2, \dots, T_p\}$ are fixed. (These numbers are meaningless for ordinary seasonal models.) Then, given these fixed parameters, several AR models of fixed order are fitted. In particular, $M = \sigma$ "auto regressive" models of order q_m , $m = 1, \dots, M$ are fitted for seasonal AR models and first-moment RAMs; $M = \sigma(p) + 1$ AR models of order q_m , $m = 1, \dots, M$ are fitted for full RAMs. The fitting is accomplished by first writing the mth AR model in vector/matrix form

$$\mathbf{x} = \mathbf{X}\mathbf{a} + \boldsymbol{\epsilon}. \tag{7}$$

Here, \mathbf{x} is a vector containing the k observations that fall into season m (seasonal AR model or first-moment RAM) or regime m (full RAM); \mathbf{a} is a vector of \tilde{p} model-coefficients; ϵ is a vector of k innovations and \mathbf{X} is the $k \times \tilde{p}$ "design matrix." The details, that is, the number \tilde{p} and the vector \mathbf{a} , are described below. The vector \mathbf{a} is estimated by minimizing the squared length

of the innovation vector. The resulting least squares estimate is given by

$$\hat{\mathbf{a}} = (\mathbf{X}^{\mathrm{T}}\mathbf{X})^{-1}\mathbf{X}^{\mathrm{T}}\mathbf{x}. \tag{8}$$

The corresponding approximate MLE estimate of the variance of the innovations is

$$\hat{\sigma}^2 = \mathbf{x}^{\mathrm{T}} (\mathbf{I} - \mathbf{X} (\mathbf{X}^{\mathrm{T}} \mathbf{X})^{-1} \mathbf{X}^{\mathrm{T}}) \mathbf{x} / k \tag{9}$$

where I denotes the identity matrix.

The goodness of the mth AR model is expressed by the Akaike information criterion (AIC) (Akaike 1973; Tong 1983)

$$AIC(q_m, d, T) = k \ln \hat{\sigma}^2 + 2\tilde{p}. \tag{10}$$

AIC rewards a close fit through the $\ln \hat{\sigma}^2$ term while imposing a penalty for excessive numbers of parameters through the $2\tilde{p}$ term. The order of the *m*th AR model is determined by the choice of q_m , $0 \le q_m \le p$, which minimizes (10).

The goodness of the entire seasonal AR model or RAM fitted in this piecewise manner with fixed delay d and threshold set \mathcal{I} is expressed as the "total AIC" (11), which is defined as the total of the AIC values summed over all models:

$$TAIC(d, \mathbb{T}) = \frac{\sum_{m=1}^{M} AIC(q_m, d, \mathbb{T})}{\sum_{m=1}^{M} k(m)}.$$
 (11)

The total AIC is divided by the total number of observations used to do the fitting to permit comparison between models using different values of delay parameter d

The ultimate model is determined by searching for the delay parameter d and threshold set \mathbb{I} , which minimizes (11). We do this via a simple brute force search procedure. Threshold sets \mathbb{I} are obtained by taking subsets of size y from a set \mathcal{T} of $k \gg y$ candidate threshold points on the real line. We simply chose \mathcal{T}

to be the collection of midpoints between all sorted realizations of the ancillary time series. Then (11) is evaluated for all combinations of subsets $\mathbb{Z} \subset \mathcal{T}$ and delay parameters d, $0 \le d \le d_{\max}$.

The details of the procedure outlined above are specified in the following three subsections. That is, the vectors \mathbf{x} and \mathbf{a} in (7), the matrix \mathbf{X} in (7) and the number \tilde{p} in (11) are described for each model m.

1) SEASONAL AR MODELS

The models in different seasons are not related to each other so that the minimization of (11) may be done for each season independently; i.e., by minimizing (10). Let $t_1, t_2, \ldots, t_{k(s)}$ be the observing times that fall into season s. Then the k(s)-vector of observations \mathbf{x} is $(X_{t_1}, X_{t_2}, \ldots, X_{t_{k(s)}})^T$, the number of coefficients $\tilde{p} = q + 1$, the vector of coefficients \mathbf{a} is $(a_0^s, a_1^s, \ldots, a_q^s)^T$ and the $k(s) \times (q + 1)$ design matrix \mathbf{x} is given by

$$\mathbf{X} = \begin{pmatrix} 1 & X_{t_1-1} & X_{t_1-2} & \cdots & X_{t_1-q} \\ 1 & X_{t_2-1} & X_{t_2-2} & \cdots & X_{t_2-q} \\ \vdots & \vdots & & \vdots & & \vdots \\ 1 & X_{t_k(s)}-1 & X_{t_k(s)}-2 & \cdots & X_{t_k(s)}-q \end{pmatrix}.$$

2) FIRST-MOMENT RAMS

As above, models in different seasons are not related to each other because the AR coefficients, except for a_o , are not affected by the ancillary time series. Again, let $t_1, t_2, \ldots, t_{k(s)}$ be the observing times that fall into season s. The k(s)-vector of observations is the same as in the seasonal AR model. However, the constant term in this model is modulated by the ancillary time series and hence a contains $\varphi + 1$ constant coefficients as well as the q dynamic coefficients. Thus,

$$\mathbf{a} = (a_0^{s,1}, a_0^{s,2}, \cdots, a_0^{s,y+1}, a_1^s, \cdots, a_q^s)^T.$$

The total number of coefficients is $\tilde{p} = q + y + 1$. The corresponding $k(s) \times \tilde{p}$ design matrix **X** has the form

$$\mathbf{X} = \begin{pmatrix} \delta[1, y(t_1)] & \cdots & \delta[\ \mathscr{J}+1, y(t_1)] & X_{t_1-1} & X_{t_1-2} & \cdots & X_{t_1-q} \\ \delta[1, y(t_2)] & \cdots & \delta[\ \mathscr{J}+1, y(t_2)] & X_{t_2-1} & X_{t_2-2} & \cdots & X_{t_2-q} \\ \vdots & & \vdots & & \vdots & & \vdots \\ \delta[1, y(t_k)] & \cdots & \delta[\ \mathscr{J}+1, y(t_k)] & X_{t_k-1} & X_{t_k-2} & \cdots & X_{t_k-q} \end{pmatrix}$$

where k represents k(s). The indicator function $\delta(j, y(t))$, defined by

$$\delta(j, y(t)) = \begin{cases} 1 & \text{if } j = y(t) \\ 0 & \text{otherwise} \end{cases}$$

identifies the region that is occupied by the ancillary time series at time t, and hence the constant term that is used in the model at that time.

3) THE FULL RAM

For a full RAM an individual AR model is fitted for each combination of season s and region j. Let $t_1, t_2, \ldots, t_{k(s,j)}$ be the collection of times t > q when X_t is in regime (s, j). The k(s, j)-vector of observations is $\mathbf{x} = (X_{t_1}, X_{t_2}, \cdots, X_{t_{k(s,j)}})^{\mathrm{T}}$, the number of AR model coefficients is $\tilde{p} = (q + 1)$, the coefficient vector is **a**

= $(a_0^{s,j}, a_1^{s,j}, \dots, a_q^{s,j})^T$ and the $k(s,j) \times (q+1)$ design matrix is

$$\mathbf{X} = \begin{pmatrix} 1 & X_{t_1-1} & X_{t_1-2} & \cdots & X_{t_1-q} \\ 1 & X_{t_2-1} & X_{t_2-2} & \cdots & X_{t_2-q} \\ \vdots & & & & \\ 1 & X_{t_{k(s,j)}-1} & X_{t_{k(s,j)}-2} & \cdots & X_{t_{k(s,j)}-q} \end{pmatrix}.$$

b. Model diagnostics

While the procedure described above is "objective," considerable vigilance is required to ensure that the fitted RAM is a reasonable one. Tong (1983) contains a large amount of material and wisdom regarding model diagnostics. We will only describe some of the things which we found useful and relevant in our particular applications, and we will also add some of our own advice to Tong's advice.

The most important advice is, of course, to ask whether the fitted model is physically reasonable. In the case that the fitted model is a RAM, one should also ensure that each regime (s, j) of the chosen model contains enough observations to make reliable parameter estimates. In all cases, one should ensure that the innovations behave as Gaussian white noise. We will discuss this in some detail below.

The full model permits innovations to have differing variance in each of its M component AR models. Thus, when a seasonal AR model or RAM fits well, the time series of estimated innovations $\{\epsilon_t; t = 1, \dots, n\}$ will be approximately stationary with respect to the mean and approximately serially uncorrelated. However, the variance will be time dependent. Therefore, estimated innovations within each component model must be normalized by the corresponding estimate of the innovation standard error (9) to apply standard statistical tests.

As a first step, the normalized residuals should be tested for Gaussian characteristics. Non-Gaussian behavior, as evidenced by skewed, plato-kurtic or bimodal normalized residuals, is an indication that the fitted model does not adequately explain the behavior of the modeled time series. It is thus useful to look at a relative frequency histogram and probability plot of the normalized residuals and to estimate their skewness and kurtosis.

Second, the normalized estimated innovations should be studied to see if they are serially uncorrelated. A very effective way in which to do this is to simply plot the autocorrelation function of the normalized residuals. In particular, checks should be made to determine whether autocorrelations are unusually large at small lags, at seasonal lags, or more often than would be expected. An objective test of the null hypothesis that the normalized residuals behave as a white noise process can be made with a "Portmanteau" lack of fit test (Box and Jenkins 1976; Ljung and Box 1978).

Given that the chosen RAM fits reasonably well, it

may be examined in more detail. Asymptotic confidence intervals can be constructed for the estimated AR parameters of each component AR model by using the fact that $\hat{\mathbf{a}}$ is asymptotically a multivariate Gaussian random vector with mean \mathbf{a} and covariance matrix $\sigma^2(\mathbf{X}^T\mathbf{X})^{-1}$ where \mathbf{X} is the design matrix appropriate to the component model and σ^2 is the corresponding innovation variance (Cox and Hinkley 1974). The covariance matrix is estimated by $\hat{\sigma}^2(\mathbf{X}^T\mathbf{X})^{-1}$ where $\hat{\sigma}^2$ is the appropriate estimate of the innovation variance.

Estimates of the conditional j-means, j-variances, and j-autocovariance functions can be derived from (3), (5), and (6). Plots of these quantities reveal useful information about the stochastic characteristics of the modeled time series when it is in each of the particular states. One can also display the behavior of forecasts of X_{t+1} , X_{t+2} , ... given specified values at previous times conditional upon y(t) = j for all t and the absence of noise. In this case "conditional j-forecasts" of X_{t+1} and X_{t+2} are given by

$$E(X_{t+1}|X_t, X_{t-1}, \cdots, X_{t+1-p}; j)$$

$$= a_0^{s(t+1),j} + a_1^{s(t+1),j} X_t + a_2^{s(t+1),j} X_{t-1} + \cdots + a_p^{s(t+1),j} X_{t+1-p}$$

$$E(X_{t+2}|X_t, X_{t-1}, \cdots, X_{t+1-p}; j) = a_0^{s(t+2),j} + a_1^{s(t+2),j} E(X_{t+1}|X_t, X_{t-1}, \cdots, X_{t+1-p}; j) + a_2^{s(t+2),j} X_{t-1} + \cdots + a_p^{s(t+2),j} X_{t+1-p}.$$
 (12)

Subsequent conditional *j*-forecasts are computed in like manner. The simplicity with which it is possible to construct forecast functions is one of the great advantages of RAMs as compared to other nonlinear time series models.

An important question is whether the fitted model is "better" than another model lower down the hierarchy of models described earlier; that is, does the ancillary time series contain useful information about state dependent variations in the annual cycle, and does it also contain additional useful information about state dependent variations in the second-order stochastic properties of the analyzed time series X_t ?

Such questions can be addressed using standard analysis of variance techniques by computing statistics of the form

$$\mathcal{F} = \frac{SSR_0 - SSR_1}{ndf_0 - ndf_1} \left(\frac{SSR_1}{ndf_1}\right)^{-1}$$
 (13)

where SSR_0 and SSR_1 are sums of squared estimated innovations for two models in the RAM hierarchy. The corresponding degrees of freedom, computed as the number of observations minus the number of estimated parameters, are given by ndf_0 and ndf_1 , respectively. Subscript "0" is used to indicate the simpler of the two models. Above, $\mathcal F$ measures the reduction of innovation variance that is brought about by using

a model of greater complexity, higher in the RAM hierarchy. Here, \mathcal{F} has a Fisher's F-distribution with (ndf₀ – ndf₁) and ndf₁ degrees of freedom under the null hypothesis that the simpler model adequately explains the observed variation of the fitted time series (Priestly 1981, sect. 5.4.4). In view of the fact that even the simplest model in the RAM hierarchy incorporates innovations with variance that are allowed to change cyclically, it makes little sense to compute \mathcal{F} for all seasons combined. We thus computed \mathcal{F} ratios for each season separately.

4. Application to the SO index

Our approach to modeling the behavior of the SO will be to begin with the simplest model in the RAM hierarchy that can accommodate seasonal variation in its first- and second-order characteristics, and then progress to the most complex RAM that we feel to be appropriate and can be estimated from the available data. At each step we will ask whether the additional complexity has led to a significant reduction in innovation variance and hence to a better description of SO characteristics.

When analyzing the nonstationary behavior of the SO by means of a full or first-moment RAM one has to specify:

- the time series X_t that is to be used as an SO index,
- the maximum AR order p,
- the seasonal indicator S_t ,
- an ancillary time series Y_t that is expected to control the SO,
 - the number of thresholds ψ , and
 - the maximum delay d_{max} .

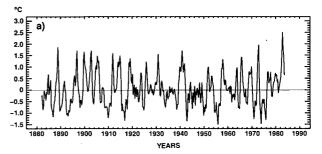
In both cases the delay parameter and the threshold set are parameters that must be estimated. In the case of the full RAM, (p+1) AR-parameters must be estimated in each regime with the implication that a total of s(y+1)(p+1)+y+1 free parameters must be specified. In the case of the 1st-moment RAM, (y+1) "constant" coefficients must be estimated for each season together with p dynamic AR coefficients. It is therefore necessary to estimate s(y+1+p)+y+1 parameters to fit the first-moment RAM. In the case of a seasonal AR model, there are only s(p+1) free parameters to estimate.

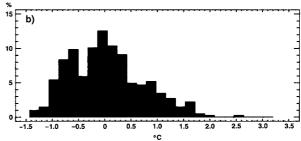
a. The SO time series X_t and the numbers p, y and d_{max}

Wright (1984) showed that there are several atmospheric and oceanic indices that carry nearly equivalent information about the SO. One of these indices is a spatial average of monthly mean sea surface temperature (SST) anomalies in the central and eastern equatorial Pacific. Wright compiled a "homogenized" SST

index of this form that is almost 150 years in length and is virtually uninterrupted during the past century. In this paper we use the 1882–1983 segment of Wright's SST time series (Fig. 1a) as an SO index. We refer to the index as the SO_{SST} index to clearly distinguish it from the usual SOI which is based on the normalized Darwin minus Tahiti SLP difference.

While there is no doubt that this long segment of Wright's SST index contains many inhomogeneities caused by data sparsity and changes in SST measurement technology, there also seems to be little doubt that this index represents the SO signal reasonably well over the entire 102-year interval. Both Darwin sea level pressure (SLP) and the Darwin minus Tahiti pressure difference (SOI) are commonly used SO indices, and Wright (1984) shows that they contain virtually the same information in the 60–20 month frequency band and on the annual time scale. He then shows that the coherence between the SO_{SST} and SLP indices in the





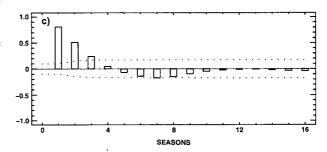


FIG. 1. Area-averaged tropical Pacific SST used as index of the state of the Southern Oscillation as compiled and homogenized by Wright (1984). (a) Time series of the SO_{SST} index. Units °C. (b) Frequency histogram of the time series shown in (a). The distribution significantly skewed. (c) Autocorrelation function of the SO_{SST} index. Lags in numbers of season. The negative autocorrelation at lag 7 seasons is significantly nonzero.

60–20 month frequency band during the period 1950–79 (when both indices are thought to be relatively free of inhomogeneities) was 0.89. The coherence in this frequency band drops to 0.75 for the 1890–1979 interval, which covers most of the 102-year interval we have selected for study. Similarly, the correlation between the annual means of the indices is 0.91 in the 1950–79 interval, and is 0.81 in the 1890–1979 interval. Thus, while the presence of inhomogeneities in SO_{SST} will introduce some uncertainty into the RAM modeling process, we are reasonably convinced that the SO_{SST} index captures enough of the SO signal during the 102-year interval we have selected to study variations in its stochastic behavior that may be related to other indicators of the state of the SO.

A frequency histogram of the SO_{SST} index is displayed in Fig. 1b. It is somewhat skewed with a wide tail in the direction of warm events. Its skewness coefficient of 0.45 is significantly greater than zero (if the effects of the annual cycle and serial correlation are ignored when making the assessment of significance). The kurtosis, -0.12, is not significantly different from that which would be expected for Gaussian samples.

To a first order of approximation, the SO_{SST} index has stochastic characteristics similar to that of an autoregressive process of order 2 (Fig. 1c). Index values at a 7-season lag are significantly negatively correlated with the current value. This is an indication of the Southern Oscillation's well-known tendency to vary on a quasi-biennial time scale (van Loon and Shea 1985, 1987). The plotted function displays average second-order stochastic characteristics for all seasons. As will be shown below and has been pointed out earlier by Chu and Katz (1985 and 1987), there is in fact considerable seasonal variation in the stochastic structure of the SO.

The SO_{SST} index varies on at least two time scales. The one we are interested in is the time scale of one year or more. The other time scale, of order one to two months, reflects the 30–60 day wave (e.g., Weickmann et al. 1985). To remove the high frequency variations we calculated three month seasonal means and set $\beta = 4$. Since a SO event (an extreme of the SO as characterized by a large excursion of the SO_{SST} or other SO index from its mean which is sustained over several seasons) usually extends from early summer to the spring of the next year (Wright 1977), a reasonable choice for the "seasons" is FMA (February, March, April), MJJ, ASO, and NDJ.

We anticipated that there might be three regions with distinct stochastic behavior, two being representative of warm or cold events and one of near normal conditions and thus set $\psi = 2$.

Chu and Katz (1985) analyzed time series of Darwin minus Tahiti SLP and inferred an AR model of order 3 using the Bayesian information criterion. However, to limit the number of parameters that must be esti-

mated from a limited amount of data, and in view of Fig. 1c, we have compromised somewhat by setting p = 2 and therefore limiting the number of free parameters to 39. We felt this to be the maximum number of parameters that could be entertained with a sample of 407 seasons. Corresponding first-moment RAMs and seasonal AR models have 23 and 12 free parameters, respectively.

We allowed the delay parameter to vary between 0 and 8, i.e., $d_{\text{max}} = 8$ seasons, thereby providing sufficient flexibility for the model to correctly align the SO_{SST} and ancillary time series, and to uncover any short term (up to four seasons) lagged relationships.

b. The seasonal index S_t and the ancillary time series Y_t

The seasonal indicator is

$$s(t) = \begin{cases} 1 & \text{if } t \in \{\text{February, March, April}\} \\ 2 & \text{if } t \in \{\text{May, June, July}\} \\ 3 & \text{if } t \in \{\text{August, September, October}\} \\ 4 & \text{if } t \in \{\text{November, December, January}\}. \end{cases}$$

$$(15)$$

To identify an appropriate ancillary time series Y_t , two different SO scenarios were exploited: the "South Pacific convergence zone" hypothesis (SPCZ) and the "Indian Monsoon" hypothesis (IM). The corresponding ancillary time series are called the SPCZ- and IMindex; and the resulting RAMs are called the SPCZ-model and the IM-model.

1) THE SPCZ HYPOTHESIS

The SPCZ hypothesis, suggested by van Loon and Shea (1985, 1987), argues that the intensity of the meridional flow in the SPCZ region, i.e., in the southwest Pacific during northern summer, is instrumental in triggering SO events. A good indicator of the strength of this meridional flow is the SLP difference between Adelaide and Rapa. Unfortunately, observations from Rapa are available only from 1952, whereas a reliable 1882-1983 Adelaide SLP time series can be obtained from the World Mean Station Climatology [available from the National Center for Atmospheric Research (NCAR)]. However, correlation of the Adelaide SLP time series with the Adelaide minus Rapa SLP time series reveals that the relevant information on the time scale we are interested in is contained in Adelaide seasonal mean SLP alone. This is clearly demonstrated by Fig. 2, which displays the Adelaide SLP and the Adelaide minus Rapa time series. We therefore used Adelaide SLP to construct the controlling time series in our modeling experiments.

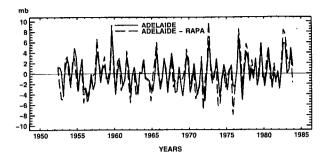


FIG. 2. NDJ mean Adelaide SLP (solid line) and NDJ Adelaide minus Rapa SLP (dashed line).

The few missing values in the Adelaide SLP time series were replaced by simple linear interpolation. Seasonal means, beginning with FMA 1882, were computed and the resulting series was seasonally adjusted by subtracting the annual cycle of the mean, and scaling by the annual cycle of the standard deviation. Low frequency variations in the mean and the variance on time scales beyond a decade were removed by subtracting the 20-year running mean from the time series and dividing by a 20-year moving estimate of the standard deviation. For the purposes of filtering, the time series of running means and moving standard errors was extrapolated at both ends so that the resulting time series of differences was again the same length as the original seasonally adjusted time series.

Finally, following the advice of van Loon and Shea (1985, 1987), we tried to enhance the signal in the controlling time series by calculating four and five season differences of the adjusted Adelaide SLP series. The physical rational for doing so is the tendency of the Southern Oscillation to oscillate on a quasi-biennial time scale (Fig. 1c). We found the five season differences to be most useful in our subsequent modeling.

A SLP tendency time series Z^{MJJ} was constructed from the adjusted and differenced Adelaide SLP series as follows:

 Z^{MJJ} = adjusted SLP in MJJ₀

$$-$$
 adjusted SLP in FMA₋₁ (16)

where subscripts "0" and "-1" indicate the present and previous calendar year. SLP tendency series Z^{ASO} , Z^{NDJ} , and Z^{FMA} were computed analogously. Four candidate ancillary time series Y_t^{MJJ} , Y_t^{ASO} , Y_t^{NDJ} , and Y_t^{FMA} were defined as piecewise constant versions of the tendency series. For example, Y_t^{MJJ} is defined by

$$Y_t^{\text{MJJ}} = Z_0^{\text{MJJ}}$$
 for

$$t = \text{FMA}_0$$
, MJJ₀, ASO₀ and NDJ₁. (17)

First-moment and full RAMs were fitted with each of the resulting ancillary time series.

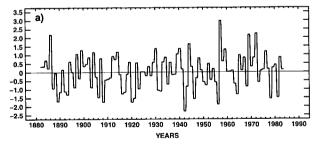
The fitted collection of RAMs were scrutinized along the lines sketched in the preceding section. First, fitted models were dismissed out of hand if the choice of thresholds resulted in regimes containing only a small number of observations. Second, the fitted model should have low normalized TAIC. This scrutiny led us to a tentative model using the $[MJJ_0 - FMA_{-1}]$ SLP index as the ancillary "SPCZ index." This choice of the controlling SPCZ index, which was made *objectively* on the basis of the normalized TAIC criterion, is consistent with the finding of van Loon and Shea (1985).

The SPCZ index and its frequency histogram are displayed in Fig. 3. The index's sample skewness (0.12) and kurtosis (-0.07) are not significantly nonzero. The autocorrelation function (not shown) has irregular characteristics, which are attributable to the long sampling interval of one year (the index is constant within years).

Notice that there is a fairly strong visual coherence between the SO_{SST} and the SPCZ indices that is confirmed by the cross-correlation functions displayed in Fig. 4. They show that MJJ, ASO and NDJ SO_{SST} values are positively correlated (≈ 0.4) with the SPCZ index at lag 0. The greatest, but relatively small, correlation in FMA occurs at lag 1 with the SO_{SST} index lagging the SPCZ index.

2) THE IM HYPOTHESIS

The IM hypothesis, suggested by Barnett et al. (1989) among others, argues that the strength of the Indian monsoon may be a precursor of extremes of the SO. An index of the strength of the Indian summer monsoon is the "all India" mean JJAS rainfall (Mooley



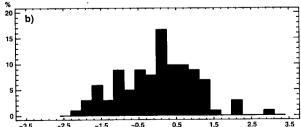


FIG. 3. The SPCZ index "Adjusted [MJJ₀ - FMA₋₁] Adelaide SLP." (a) Time series. Normalized units. (b) Frequency histogram. The data do not contradict the null hypothesis of a normal distribution.

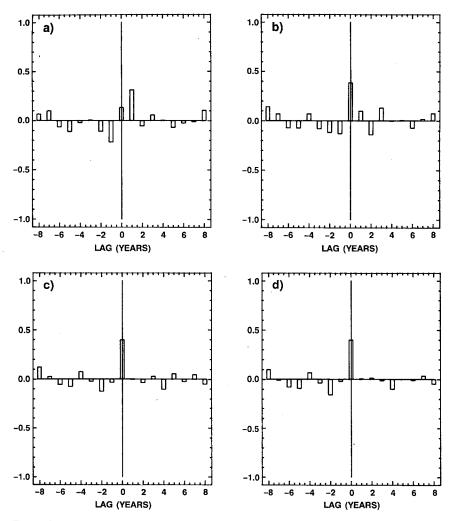


Fig. 4. Cross correlation of the SO_{SST} index and the SPCZ index stratified according to season. (a) SO_{SST} index in FMA, (b) JJA, (c) SON, and (d) NDJ. Lag in years.

and Parthasarathy 1984; Shukla 1987; Elliot and Angell 1987). This series is used as an "IM index" Y_t by fixing the time series to the JJAS value within calendar years. The IM index and its frequency histogram are displayed in Fig. 5. The index's sample skewness (-0.52) is significantly different (at the 5% level) from that expected for a Gaussian sample. Its kurtosis (-0.08) does not deviate significantly from zero.

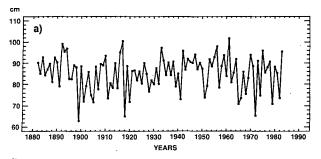
As with the SPCZ index, variations of the IM index are fairly coherent with the variations of the SO_{SST} index. This is substantiated by the cross-correlation functions shown in Fig. 6. MJJ, ASO and NDJ SSTs are negatively correlated (≈ -0.6) with the IM index at lag 0. The greatest correlation in FMA occurs at lag 1, with the SO_{SST} index following the IM index. The latter correlation is comparable to the ASO and NDJ values. These findings are in general agreement with those of Elliot and Angell (1987).

5. The fitted models

a. The ordinary seasonal AR model

We fitted an ordinary seasonal AR model to the SO_{SST} index to provide a baseline of comparison for subsequent RAMs. Details of the fitted seasonal AR model are given in Table 1. Note that our results are in general accord with those described by Chu and Katz (1985) even though the same SO index is not used in both studies and the definition of the seasons is somewhat different. In particular, both models show that the SO is least persistent in Northern Hemisphere spring. A major difference is that our model tends to amplify SO_{SST} anomalies in NDJ.

The annual march of the mean and standard deviation of the SO_{SST} index that is derived from the fitted seasonal AR model is shown in Fig. 7. The seasonal means should be and are practically zero because the



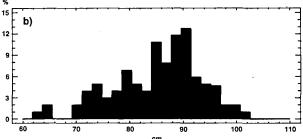


FIG. 5. The IM index "all India summer (JJAS) rainfall." (a) Time series. Units: cm. (b) Frequency histogram. The data do not contradict the null hypothesis of a normal distribution.

SO_{SST} index is derived from SST anomalies. The standard deviation exhibits a marked annual variation with a maximum in FMA and a minimum in ASO. This behavior reflects the earlier mentioned tendency of the SO to be highly persistent in its "peak" phase but highly variable at the end of an event.

b. The SPCZ models

1) THE FIRST-MOMENT SPCZ-RAM

The parameters of the fitted model are given in Table 2. The estimated innovations do not contradict the hypothesis that they behave as stationary white noise after normalization. Their standard deviations are approximately one-half that of the SO index.

The estimated thresholds are -1.12 and 0.84; i.e., the model is in region j = 1 if $Y_t \le -1.12$, j = 2 if $-1.12 < Y_t \le 0.84$, and j = 3 if $Y_t > 0.84$. These values were estimated by using the search procedure outlined in section 3a. Negative values of the ancillary time series refer to negative 15-month SLP change at Adelaide. or according to Fig. 2, a negative change in Adelaide minus Rapa pressure; i.e., strengthening southerly winds in the southwest Pacific. Similarly, positive values are indicative of strengthening northerly winds. Therefore, regime j = 1 is called "southerly," j = 2"normal," and j = 3 "northerly." The delay parameter was chosen to be one season (d = 1) suggesting that the 15-month SLP change at Adelaide between MJJ of the current year and FMA of the previous year affects the course of the annual cycle of the SO_{SST} index between the current MJJ and the next FMA.

The annual cycle of the conditional *j*-means of the first-moment SPCZ RAM are displayed in Fig. 7a. The normal means are close to zero in all seasons, the southerly means fluctuate around -0.55° C and the northerly means vary between $+0.6^{\circ}$ C and $+0.9^{\circ}$ C; that is, the southerly regimes tend to be connected with below normal tropical Pacific SST (cold event), and the northerly regimes with above normal SST (warm event). This is the scenario that was suggested by van Loon and Shea (1985). In the following we identify southerly with "cold," and northerly with "warm." Interestingly, the mean warm events are described as being larger than mean cold events.

The annual march of the j-standard deviations (Fig. 7b), which in the case of first-moment models does not depend on the regime j, is similar to that of the seasonal AR model with respect to the pattern, but the annual mean of the standard deviation is reduced by about 0.1°C. This reduction reflects the RAM's success in using the state of the ancillary time series to modulate the mean state and hence in reducing the magnitude of anomalies from the annual cycle.

The first-moment SPCZ RAM is superior to the ordinary seasonal AR model; the F-statistic in FMA, MJJ, and ASO is sufficiently large to reject the null hypothesis that there is no additional information in the ancillary time series, i.e., in the SPCZ index (risks: 2.0% in FMA, ≪0.1% in MJJ, and 3.9% in ASO).

The results of the \mathcal{F} -tests are consistent with the details of the first-moment RAM. There is little difference between the dynamic components, a_1 and a_2 , of the seasonal AR model (Table 1) and the first-moment RAM (Table 2). However, the seasonal AR model incorporates only a weak annual cycle, whereas the first-moment RAM incorporates an annual cycle that is strongly modulated by the SPCZ index as is evident from the difference between the "warm/northerly," "normal," and "cold/southerly" annual cycles displayed in Fig. 7a.

Insight into the behavior of the dynamic part of the first-moment RAM can be obtained by considering the conditional *j*-forecasts (12) made from a 1°C anomaly of one season duration. These forecasts are displayed in Fig. 8 as anomalies from the *j*-mean. Note that because the dynamic part of the first-moment RAM is not affected by the ancillary time series, there is only one forecast of anomalies from the conditional mean per season. The figure shows that an SO_{SST} anomaly from the conditional mean that appears in

- NDJ decays over the course of four to five seasons in the absence of any forcing.
- FMA persists until the following NDJ before decaying exponentially over the course of two seasons.
- MJJ is amplified for two seasons so that it attains about 150% of its original magnitude in NDJ. It then quickly decays over the course of two to three seasons.

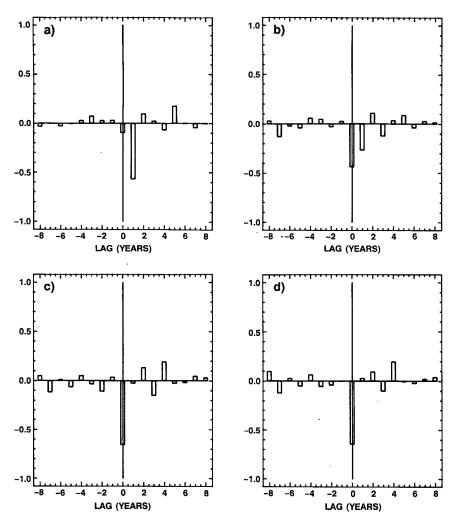


FIG. 6. Cross correlation of the SO_{SST} index and the IM index stratified according to season.
(a) SO_{SST} index in FMA, (b) JJA, (c) SON, and (d) NDJ. Lag in years.

• ASO is slightly amplified for one season and is then again quickly damped.

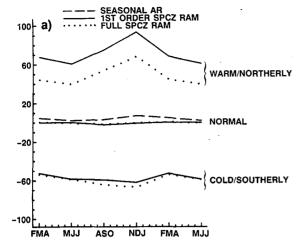
2) THE FULL SPCZ-RAM

The estimated parameters of the full SPCZ model are also given in Table 2. The estimated innovations do not contradict the hypothesis that the normalized innovations come from a white noise process. The same delay and threshold parameters were chosen as for the first-moment RAM. We will see that the most important feature of RAMs for modeling the SO_{SST} index is their ability to modulate the annual cycle. It is, thus, not unreasonable that there was not a refinement in the estimated delay and threshold parameters as they were estimated from a discrete set of possible candidates.

The annual cycle of the conditional *j*-means and *j*-standard deviations of the full RAM are displayed in

Fig. 7. The normal-means are close to zero in all seasons, the cold/southerly annual cycle varies near -0.6 °C, while the warm/northerly annual cycle varies between +0.4 °C and 0.7 °C. The j-standard deviations exhibit a marked seasonal variation. Also, variability in the warm/northerly regime is substantially larger than in the cold/southerly regime, at least in NDJ and FMA.

It is not clear whether the full SPCZ-RAM is superior to the first-moment model described above or whether it is over specified. First, there is only a small increase in explained variance beyond the first-moment model. Second, the null hypothesis that the ancillary time series does not modulate the model's dynamic components, a_1 , a_2 , and ϵ , can be rejected only in MJJ and ASO (risk: $\ll 0.1\%$ and 4.9%, respectively). This is consistent with differences that can be noted in Table 2. The AR models fitted in FMA and NDJ are very similar to those fitted for the first-moment RAM. But, the models



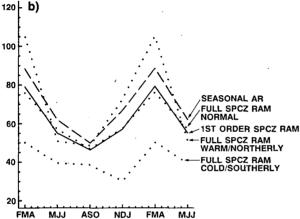


FIG. 7. Annual cycle of (a) *j*-means, and (b) *j*-standard deviations. Dashed line—seasonal AR model, solid lines—1st-moment SPCZ RAM, dotted lines—full SPCZ RAM. Note that FMA and MJJ are repeated. Also note that "normal" first-moment and full RAM curves overlay each other in panel (a) Units: 10⁻² °C.

in the cold region in MJJ and ASO exhibit strikingly different stochastic behavior. The differences suggest that disturbances are more persistent during a warm event than during a cold event in the sense that anomalies about the modulated annual cycle are damped much more vigorously in MJJ and ASO during cold regimes than during warm regimes. This may in part explain the systematic difference of the *j*-standard deviations mentioned above.

The difference in stochastic behavior between warm and cold regimes can also be seen in the conditional *j*-forecasts (Fig. 9). Forecasts made from SO_{SST} anomalies that appear in NDJ and ASO evolve in a manner similar to that of the corresponding forecasts derived from the first-moment model. One minor difference from the first-moment model in the warm case is that a positive anomaly is followed by a small negative anomaly in NDJ of the subsequent year. Like the first-

moment model, an SO_{SST} anomaly that appears in FMA or MJJ decays in the model in approximately 6–8 seasons. However, unlike the first-moment model, there are considerable differences between the warm, normal, and cold region forecasts.

The full SPCZ model is not nearly as well specified as the first-moment model because smaller numbers of observations are used to estimate individual parameters. Consequently, there is a loss of precision in estimated quantities such as the conditional *j*-means. Unfortunately, this loss can not be quantified and one has to suspect that at least part of the details of the SPCZ RAM merely mirror sampling properties.

c. The IM models

Three-region IM models similar to the SPCZ models were fitted to the data with disappointing success. In both cases the threshold search procedure led to models that were in essence 2-region RAMs in that two of the three defined regions contained all but a few of the observed IM precipitation amounts. We therefore decided to fit 2-region first-moment and full IM RAMs.

1) THE FIRST-MOMENT IM-RAM

The parameters of the fitted 2-region first-moment IM model are given in Table 3. As with the SPCZ model, this model explains a large proportion of the variance of the SO_{SST} index, and the Portmanteau test does not contradict the null hypothesis that the normalized residuals come from a white noise process. The single threshold is located at 85.7 cm of IM rainfall, very near the 102-year mean IM rainfall of 85.2 cm. We refer to the two regions defined by the threshold as the "dry" and "wet" regions. The IM hypothesis predicts that a wet monsoon will be associated with a "cold" SO event and vice versa.

The delay parameter was chosen to be one season. Therefore, the form of the annual cycle of SO_{SST} be-

TABLE 1. Ordinary seasonal AR model fitted to the 1882–1983 SO_{SST} index with maximum AR order 2 permitted in any season. Units (for a_0 and σ_c): 10^{-2} °C. Estimates of a_1 and a_2 are dimensionless. Quantities in parentheses indicate asymptotic standard errors of the corresponding parameter estimates.

| | | Std dev | | |
|--------|--------|---------|-----------------------|---------------------------|
| Season | a_0 | a_1 | <i>a</i> ₂ | $\hat{\sigma}_{\epsilon}$ |
| FMA | 0.39 | 0.571 | • | 33.2 |
| | (3.31) | (0.037) | | |
| MJJ | -0.17 | 1.032 | -0.368 | 37.4 |
| | (3.73) | (0.112) | (0.076) | |
| ASO | 2.55 | 1.436 | -0.471 | 36.2 |
| | (3.59) | (0.087) | (0.077) | |
| NDJ | 3.56 | 1.172 | , , | 27.1 |
| | (2.70) | (0.037) | | |

TABLE 2. First-moment and full SPCZ-RAM with 3_{δ} regimes and maximum AR order 2 as fitted to the SO_{SST} index. The change in Adelaide SLP between FMA of the previous year and MJJ of the current year is used as the ancillary time series. The thresholds, T_1 and T_2 , are located at -1.12 and +0.84 of the normalized SPCZ index. Units (for a_0 and σ_s): 10^{-2} °C. Estimates of a_1 and a_2 are dimensionless. Quantities in parentheses indicate asymptotic standard errors of the corresponding parameter estimates.

| Season Region | | First-moment SPCZ RAM | | | | | Full SPCZ RAM | | | | |
|---------------|--------|-----------------------|-----------------------|-------------------|---------|-------------------|-----------------------|-------------------|---------------------------|--|--|
| | | | Parameters | | G(4 4 | Parameters | | | Std dev | | |
| | Region | <i>a</i> ₀ | <i>a</i> ₁ | a ₂ | Std dev | a_0 | <i>a</i> ₁ | a ₂ | $\hat{\sigma}_{\epsilon}$ | | |
| FMA | 1 | -17.02 (7.98) | 0.572 (0.035) | | 31.9 | -16.87 (8.09) | 0.554 (0.072) | | 32.2 | | |
| | 2 | 0.56 (3.93) | 0.572 (0.035) | | 31.9 | 0.60 (3.57) | 0.609 (0.039) | | 28.7 | | |
| | 3 | 14.40 (7.32) | 0.572 (0.035) | | 31.9 | 13.39 (8.72) | 0.455 (0.103) | | 39.0 | | |
| MJJ | 1 | -29.68 (8.66) | 0.894 (0.106) | -0.288 (0.071) | 33.9 | -38.37 (7.35) | 0.381 (0.117) | | 28.8 | | |
| | 2 | -0.65 (4.17) | 0.894 (0.106) | -0.288 (0.071) | 33.9 | -1.17 (4.45) | 0.933 (0.155) | -0.287 (0.106) | 35.8 | | |
| | 3 | 26.60 (7.92) | 0.894 (0.106) | -0.288 (0.071) | 33.9 | 24.00 (5.00) | 1.107 (0.121) | -0.496 (0.079) | 21.1 | | |
| ASO | 1 | -3.41 (9.44) | 1.344 (0.095) | -0.437 (0.077) | 35.0 | -32.78 (13.70) | 0.538 (0.241) | | 35.9 | | |
| | 2 | -1.98 (4.28) | (0.095) | -0.437 (0.077) | 35.0 | -0.64 (3.91) | 1.397 (0.103) | -0.506 (0.086) | 31.7 | | |
| | 3 | 23.44 (8.62) | 1.344 (0.095) | -0.437 (0.077) | 35.0 | 9.10 (12.37) | 1.563 (0.255) | -0.379 (0.244) | 41.7 | | |
| NDJ | 1 | 7.67 (7.18) | 1.174 (0.043) | | 27.1 | 2.21 (9.81) | 1.076 (0.108) | | 23.2 | | |
| | 2 | 2.07 (3.34) | 1.174 (0.043) | | 27.1 | 1.83 (3.58) | 1.140 (0.056) | | 28.8 | | |
| | 3 | 5.26 (6.91) | 1.174 (0.043) | | 27.1 | -0.85 (6.37) | 1.273 (0.064) | | 21.8 | | |

tween the current MJJ and the following FMA, as represented by the fitted model, is influenced by the concurrent JJAS Indian monsoon precipitation. The annual cycle of the conditional j-means and j-standard deviations of the IM model are displayed in Fig. 10. We see that the model has means of the expected sign in the wet and dry regions and thus we designated these

regions as "wet/cold" and "dry/warm." Consistent with the first-moment SPCZ RAM, the *j*-standard deviations of the first-moment IM RAM are about 0.1°C less than those of the seasonal AR model.

The comparison of parameter estimates for the first-moment IM-RAM (Table 3) and for the seasonal AR model (Table 1) reveal that the main difference be-

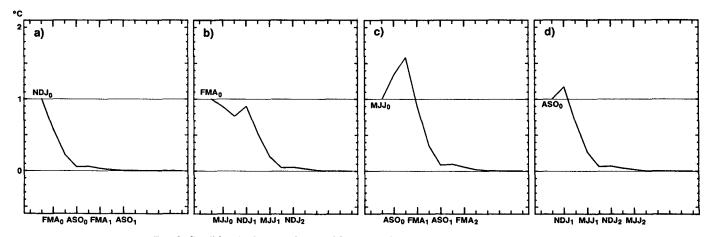


Fig. 8. Conditional j-forecasts from a 1°C anomaly imposed at season (a) NDJ, (b) FMA, (c) MJJ, and (d) ASO made with the first-moment SPCZ model.

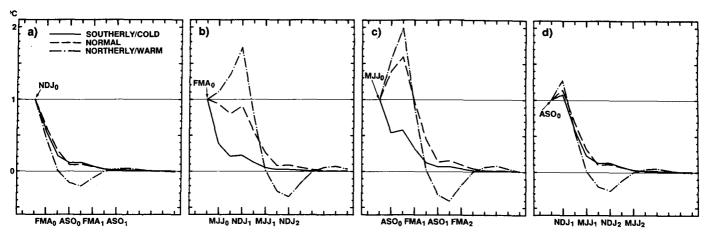


FIG. 9. Conditional j-forecasts from a 1°C anomaly imposed at season (a) NDJ, (b) FMA, (c) MJJ, and (d) ASO made with the full SPCZ model.

tween the two models is in the estimates of the constant terms in MJJ and ASO. Thus, the main difference between the two models is that the annual cycle is modulated in MJJ and ASO in the first-moment model by an indicator of whether the concurrent monsoon is weaker or stronger than normal. The \mathcal{F} -test supports this interpretation by rejecting the hypothesis that the IM-index does not modulate the annual cycle of SO_{SST} in MJJ and ASO (risk $\leq 0.1\%$ in both seasons).

2) THE FULL IM-RAM

The details of the full 2-region IM RAM will not be described. There is only a weak suggestion that a two-state index of the Indian monsoon affects the (second-order) stochastic properties of the RAM. The null hypothesis that the two-state IM index does not modulate the model's dynamic components can be rejected only in ASO with a relatively large risk (4.9%).

TABLE 3. First-moment IM-RAM with 2σ regimes and maximum AR order 2 as fitted to the SO_{SST} index. The threshold is located at 85.7 cm of IM precipitation. Units (for a_0 and σ_t): 10^{-2} °C. Estimates a_1 and a_2 are dimensionless. Quantities in parentheses indicate asymptotic standard errors of the corresponding parameter estimates.

| | | Parameters | | | | | |
|--------|-----------------|------------------|------------------|-------------------|---------|--|--|
| Season | a_0^1 | a_0^2 | a_1 | a ₂ | Std dev | | |
| FMA | -1.22 (5.91) | 1.62 (5.02) | 0.581 (0.048) | | 33.1 | | |
| MJJ | 20.17 (4.91) | -17.28 (4.49) | 0.904 (0.101) | -0.276 (0.069) | 32.7 | | |
| ASO | 25.77 (5.09) | -16.89 (4.59) | 1.164 (0.089) | -0.331 (0.071) | 31.4 | | |
| NDJ | 4.99 (4.85) | 2.37 (4.23) | 1.160 (0.050) | | 27.1 | | |

d. Classification of ENSO events

One of the by-products of our RAM analysis is a times series of SO state estimates as characterized by the sequence of regimes identified by the RAM and the labels (cold, normal, warm), which we have attached to these regimes. On the other hand, it is possible to categorize each year as being the year of a "cold," "warm," or "no" event by means of the SO_{SST} index itself. If the NDJ SO_{SST} is less (greater) than the mean minus (plus) one standard deviation of the index the year is labeled "cold" ("warm"). All other years are labeled "no" years.

A measure of association between the two classifications, by the RAM and by the SO_{SST} index alone, may be obtained by compiling a contingency table and conducting a χ^2 -test of the null hypothesis that the two partitions are unrelated. Strong association is not a measure of success because the indices used as ancillary variables are correlated with the SO_{SST} time series. On the other hand, lack of association indicates that thresholds have been chosen in such a way as to obliterate the connection between the index and SO_{SST} series.

Since the first-moment SPCZ RAM and the full SPCZ RAM operate with the same thresholds, the classifications made by the models are identical. In the contingency table derived from the SPCZ RAM (Table 4), there is a strong association between the RAM and SO_{SST} classifications of the SO as measured by the χ^2 statistic. Comparison of the column and row sums of Table 4 shows that the SPCZ-classifications have a distribution similar to that of the observed states. Examination of the cells of this table shows that warm (cold) classifications coincide with warm (cold) SO years more frequently than expected if there were no association between estimates and observations.

SPCZ-classifications also have stochastic character-

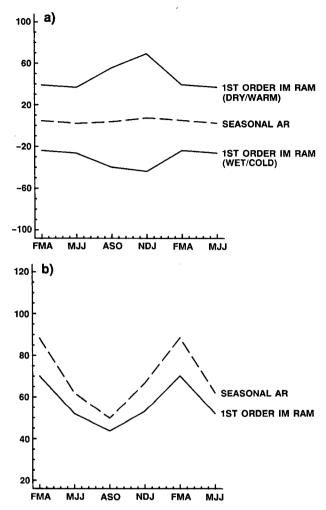


Fig. 10. Annual cycle of (a) *j*-means, and (b) *j*-standard deviations for the 2-region IM RAMs. Dashed line—seasonal AR model, solid lines—first-moment IM RAM. Note that FMA and MJJ are repeated. Units: 10⁻² °C.

istics that are similar to those of the observed states as is revealed by Table 5. Both the occupation times and transition probabilities are similar in the SO_{SST} classifications and the SPCZ-RAM classifications. A severe deviation is the cell "warm at time t, cold at time t+1," which is zero in the SPCZ RAM. A minor deviation is that when the SPCZ classification is either warm or cold there is a greater than observed probability that the classification the following year will be normal.

6. Summary

We have introduced a fairly general class of nonlinear time series models—regime-dependent autoregressive models (RAMs)—into climate research and demonstrated its usefulness by describing the Southern Os-

cillation. The basic idea of the RAM is to define different regions in the considered state space and to fit within each region a regular seasonal AR model. The choice of which region is appropriate at a given time is made by an ancillary variable.

In our example we use the models to describe the Southern Oscillation as characterized by Wright's SST index (Wright 1984). We exploited two different hypotheses of how the Southern Oscillation is maintained to obtain ancillary time series; the "South Pacific convergence zone (SPCZ)-hypothesis" of van Loon and Shea (1985, 1987) and the "Indian monsoon (IM) hypothesis" as reported by Barnett et al. (1989) and others. The SPCZ hypothesis suggests the use of Adelaide sea level pressure as the ancillary time series and the Indian monsoon rainfall appears to be a reasonable choice for the IM hypothesis.

Both SPCZ RAMs, the first-moment and the full model, turned out to be skillful descriptions of the Southern Oscillation. These models represent a significant advance over previous time series models of the SO because they incorporate a representation of the annual cycle that is dynamically modulated by an indicator of the state of the SO. The inclusion of these annual cycle terms results in a significant reduction of variance attributed to noise when compared with ordinary seasonal AR models that allow for only a fixed, nonstate dependent representation of the annual cycle. Both RAM models suggest that events during FMA are pivotal in determining the course of the SO during the subsequent year. An indication of which course has been taken is available at the end of July in the form of the 15-month Adelaide FMA₋₁ to MJJ₀ SLP tendency. The characteristics of the full SPCZ-RAM indicate that the stochastic characteristics of the SO index may also depend on the regime. In the model,

TABLE 4. Contingency table describing the association between the classification of years as cold, normal, or warm SO years using the SO_{SST} index (labeled "observed"; "cold/warm" corresponds to NDJ SST less/greater than the mean minus/plus 1 standard deviation) and the classification obtained from the SPCZ-RAM.

Quantities in parentheses are expected cell frequencies under the hypothesis that there is no association. The χ^2 -statistic for the table is 33.2 with 4 degrees of freedom (significant at the 0.01% level provided the data were not serially correlated).

| | | Observed SO state | | | В |
|--------------------|--------------------|-------------------|--------------|--------------|--------------|
| | | Cold | Normal | Warm | Row total |
| SPCZ RAM: | Cold/ Southerly | 9 (3.5) | 7 (8.5) | 0 (4.0) | 16 |
| Estimated SO State | Normal | 12 (14.2) | 41 (34.7) | 12 (16.1) | 65 |
| | Warm/ Northerly | 1 (4.4) | 6 (10.7) | 13 (4.9) | 20 |
| | Col Total | 22 | 54 | 25 | 101 |

TABLE 5. Transition probability matrix using the classifications obtained from the SOSST index and from the SPCZ-RAM (101 years).

| | State at time $t+1$ | | | | | | |
|----------------------|-------------------------|--------|------|-------------|--------|------|--|
| | SO _{SST} index | | | SPCZ RAM | | | |
| State at time t | Cold | Normal | Warm | Cold | Normal | Warm | |
| Cold | 32% | 36% | 32% | 25% | 56% | 19% | |
| Normal | 17% | 61% | 22% | 18% | 57% | 25% | |
| Warm | 21% | 54% | 25% | '0 % | 95% | 5% | |
| Mean occ. time (yrs) | 1.5 | 2.6 | 1.3 | 1.3 | 2.3 | 1.1 | |

the course of a cold event is less variable than that of a warm event. We cannot exclude, however, the possibility that this finding is due to over fitting the data. It is fair to say that the full RAMs employed in this paper stretch the limits of our ability to estimate parameters with seasonal time series that are of order 100 years in length.

The IM hypothesis leads to a less successful RAM. In contrast to the 3-region SPCZ models, the IM model operates in only two regions. The sign of the IM precipitation anomaly may be used to infer which of the two possible regimes the SO is in at a given time. The "above normal" and "below normal" regimes differ mainly with respect to the annual cycle. The extension of the first-moment model to a full RAM provided only weak evidence that the IM index affects the second moments of the SO_{SST} index. The fact that the IM hypothesis does not result in a workable 3-region model with distinct warm, cold, and intermediate regime characteristics suggests that the Indian monsoon does not contain a useful SO precursor signal. Rather, the IM RAMs describe a concurrent relationship between the SO and the Indian monsoon. This is consistent with the findings of Elliot and Angell (1987) who showed that the IM index was not correlated with prior SO variation.

An important caveat to this study is that it has been conducted very much in an exploratory vein. We have demonstrated that RAMs are able to capture stochastic features of the SO and its relationships with certain precursor time series that have previously been described in other ways. This illustrates the potential utility of this type of statistical model. However, it does not constitute an independent confirmation of the existence of these stochastic relationships because our findings are not based on independent data. Ultimate confirmation or rejection of the SPCZ- and IM-hypotheses awaits either future climate observations or the development of coupled atmosphere/ocean GCMs that can simulate the ENSO phenomenon with considerable fidelity.

This study illustrates both the advantages and the potential problems of RAMs. These models can be interpreted as discretized nonlinear differential equations. As such, this is a very broad class of models that can

be used without precise or restrictive specification of the form of the governing equations. In that sense these models are almost nonparametric and lend themselves well to an exploratory analysis of geophysical data. However, the price of this generality is the need to estimate numerous parameters. Our experience with these models shows that there is a need for a reasonable physical model of the process that can be used to constrain the class of statistical models to be considered and, thus, strengthen the inferences that are possible.

Acknowledgments. The authors gratefully acknowledge support for this work provided by the Max Planck Institut für Meteorologie to Francis Zwiers for the purpose of facilitating this collaboration.

REFERENCES

- Akaike, H., 1973: Maximum likelihood identification of Gaussian autoregressive moving average models. *Biometrika*, **60**, 255–265
- Barnett, T. P., 1983: Interaction of the monsoon and Pacific trade wind systems at interannual time scales. Part I: The equatorial zone. *Mon. Wea. Rev.*, 111, 756-773.
- —, 1984a: Interaction of the monsoon and Pacific trade wind systems at interannual time scales. Part II: The tropical band. *Mon. Wea. Rev.*, **112**, 2380–2387.
- —, 1984b: Interaction of the monsoon and Pacific trade wind systems at interannual time scales. Part III: A partial anatomy of the Southern Oscillation. Mon. Wea. Rev., 112, 2388-2400.
 —, 1985: Variations in near-global sea level pressure. J. Atmos.
- Sci., 42, 478-501.
- ----, L. Dümenil, U. Schlese, E. Roeckner and M. Latif, 1989: The effect of Eurasian snow cover on regional and global climate variations. *J. Atmos. Sci.*, 46, 661-685.
- Box, G. E. P., and G. M. Jenkins, 1976: Time Series Analysis Forecasting and Control. Revised Edition, Holden-Day, 575 pp.
- Chu, P.-S., and R. W. Katz, 1985: Modeling and forecasting the southern oscillation: a time domain approach. Mon. Wea. Rev., 113, 1876-1888.
- , and , 1987: Measures of predictability with applications to the Southern Oscillation. Mon. Wea. Rev., 115, 1542-1549.
 , and , 1989: Spectral estimation from time series models with relevance to the Southern Oscillation. J. Climate, 2, 86-
- 90. Cox, D. R., and D. V. Hinkley, 1974: *Theoretical Statistics*. Chapman Hall, 511 pp.
- Elliot, W. P., and J. K. Angell, 1987: The relationship between Indian monsoon rainfall, the Southern Oscillation, and hemispheric air and sea temperature: 1884–1984. *J. Climate Appl. Meteor.*, 26, 943–948.

- Fraedrich, K., 1988: El Niño/Southern Oscillation predictability. Mon. Wea. Rev., 116, 1001-1012.
- Gutzler, D. S., and D. E. Harrison, 1987: The structure and evolution of seasonal wind anomalies over the near equatorial eastern Indian and western Pacific Oceans. *Mon. Wea. Rev.*, 115, 169– 192.
- Hense, A., 1987: On the possible existence of a strange attractor for the Southern Oscillation. *Beitr. Phys. Atmos.*, **60**, 34-47.
- Ljung, G. M., and G. E. P. Box, 1978: On a measure of lack of fit in time series models. *Biometrika*. 65, 297-303.
- Mooley, D. A., and B. Parthasarathy, 1984: Fluctuations in All-India summer monsoon rainfall during 1871–1978. *Climate Change*, **6.** 287–302.
- Priestly, M. B., 1981: Spectral Analysis and Time Series. Academic Press, 890 pp.
- Rasmusson, E. M., and T. H. Carpenter, 1982: Variations in tropical sea surface temperature and surface wind fields associated with the Southern Oscillation/El Niño. Mon. Wea. Rev., 110, 354– 384.
- Shukla, J., 1987: Interannual variability of monsoons. *Monsoons*, J. S. Fein and P. L. Stephens, Eds., Wiley, 399-463.
- Tong, H., 1983: Threshold Models in Nonlinear Time Series Analysis. Springer-Verlag, 323 pp.

- van Loon, H., and D. J. Shea, 1985: The Southern Oscillation. Part IV: The precursors south of 15°S to the extremes of the oscillation. *Mon. Wea. Rev.*, 113, 2063–2074.
- —, and —, 1987: The Southern Oscillation. Part VI: Anomalies of sea level pressure on the Southern Hemisphere and of the Pacific sea surface temperature during the development of a warm event. *Mon. Wea. Rev.*, 115, 370–379.
- Weickmann, K. M., G. R. Lussky and J. E. Kutzbach, 1985: Intraseasonal (30–60 day) fluctuations of outgoing longwave radiation and 250 mb stream function during northern winter. *Mon. Wea. Rev.*, 113, 941–961.
- Wright, P. B., 1977: The Southern Oscillation—Patterns and Mechanisms of the Teleconnections and Persistence. Hawaii Institute of Geophysics Publ. 77-13.
- —, 1984: Relationships between indices of the Southern Oscillation.

 Mon. Wea. Rev., 112, 1913–1919.
- —, 1985: The Southern Oscillation: an ocean-atmosphere feedback system? *Bull. Amer. Meteor. Soc.*, **66**, 398-412.
- —, J. M. Wallace, T. P. Mitchell and C. Deser, 1988: Correlation structure of the El Niño/Southern Oscillation phenomenon. J. Climate, 1, 609-625.
- Xu, J., and H. V. Storch, 1990: Principal Oscillation Pattern—prediction of the state of ENSO, *J. Climate*, **3**, 1316–1329.