

Time–Frequency Variability of ENSO and Stochastic Simulations

TAHL S. KESTIN, DAVID J. KAROLY, AND JUN-ICHI YANO

CRC for Southern Hemisphere Meteorology, Monash University, Melbourne, Australia

NICOLA A. RAYNER

Hadley Centre, Meteorological Office, Bracknell, United Kingdom

(Manuscript received 8 April 1997, in final form 17 November 1997)

ABSTRACT

The time–frequency spectral structure of El Niño–Southern Oscillation (ENSO) time series holds much information about the physical dynamics of the ENSO system. The authors have analyzed changes of the spectrum with time of three ENSO indices: the conventional Southern Oscillation index (SOI), Niño3 sea surface temperatures, and a tropical Pacific rain index, over the period 1871–1995. Three methods of time–frequency analysis—windowed Fourier transform, wavelet analysis, and windowed Prony’s method—were used, and the results are in good agreement. The time–frequency spectra of all the series show strong multidecadal variations over the past century. In particular, there was reduced activity of ENSO in the 2–3-yr periodicity range during the period 1920–60, compared with both the earlier and later periods. The dominant frequencies in the spectra do not appear to be constrained to certain frequency bands, and there is no evidence that the ENSO system has fixed modes of oscillation.

The qualitative behavior of the real SOI time series has been compared with that of time series simulated by an autoregressive stochastic process of order 3 and time series created by phase-randomizing the spectral components of the SOI. The decadal variability of the amplitude and time–frequency spectra was found to be very similar between the observed and simulated SOIs. This suggests that the decadal variability of ENSO can be well simulated by a stochastic model and that stochastic forcing may be an important component of ENSO dynamics.

1. Introduction

Understanding the nature of the variability of our climate is emerging as the primary focus of many climatological studies. The El Niño–Southern Oscillation (ENSO) is a particularly important component of climate variability, second in magnitude only to the annual cycle. It refers to coherent fluctuations of sea surface temperatures, pressure anomalies, and zonal winds in the basins of the tropical Indian and Pacific Oceans. We have used time–frequency analysis methods to investigate and describe interannual and interdecadal variability of ENSO series over the past 120 yr. Several indicators of ENSO have been used to confirm the reliability of the different series. This approach allows us to infer some dynamical characteristics of the ENSO system.

The power spectra of long ENSO-related time series (>40 yr) have a broad spectral peak corresponding to

periods of about 2–8 yr (Rasmusson and Carpenter 1982). In many cases, the broad peak of the spectra appears to be subdivided into a 2–3-yr quasi-biennial (QB) peak and a 3–8-yr low-frequency (LF) peak (Barnett 1991). Barnett (1991) and Allan et al. (1995) found that the variability associated with each band has changed over the past century. Some recent theoretical and analytical work views these bands as fundamental modes of fluctuation of the ENSO system, whose interference and interaction with each other and with the annual cycle produces the variability in ENSO fluctuations (Barnett 1991; Allan et al. 1995; Brassington 1997). The LF mode appears to be energetically more important.

Recently, Wang and Wang (1996) used wavelet and waveform analysis to closely investigate how the spectrum of ENSO series has changed over the past century. They found that there have been large changes in the dominant frequencies on the scale of decades, although these frequencies appeared to be confined to constant values for periods of up to 20 yr. Wang and Wang concluded that these results support the view that ENSO variability is produced by the interaction of several fundamental modes.

Corresponding author address: Prof. David J. Karoly, CRC for Southern Hemisphere Meteorology, Monash University, Wellington Rd., Clayton, Victoria 3168, Australia.
E-mail: djk@vortex.shm.monash.edu.au

In this paper we examine the changes in the spectrum over time of three ENSO series: the Southern Oscillation index (SOI), Niño3 SST anomalies, and a tropical Pacific rainfall index (section 2). We use three time–frequency analysis methods—windowed Fourier transforms, wavelet analysis, and the windowed Prony method—to calculate the spectra (section 3). The spectral patterns produced are in good agreement between the different time–frequency analysis methods and between the different data series. They support the existence of a single dominant mode with a continuously changing frequency in the 2–10-yr range, rather than the occurrence of multiple modes with constant frequencies. Unlike Wang and Wang (1996), we find good agreement between the time–frequency variability of the SST index and the SOI before 1950.

To assess the significance (both physical and statistical) of the variability of the ENSO spectra, two approaches have been used to simulate ENSO time series. In the first, the SOI time series was simulated by an autoregressive stochastic process of order 3. In the second approach, the phase of the individual spectral components of the observed SOI was randomized and recombined to form a new time series. This phase-randomized series has an identical power spectrum to the observed SOI but has different phase relationships between the spectral components. The decadal variability of the amplitude and time–frequency spectra of the observed SOI was compared with those of the simulated time series. This allows us to infer the significance of the observed decadal variations of the SOI spectrum.

2. Data

The ENSO phenomenon has three main aspects: an atmospheric component, an oceanic component, and teleconnection patterns. To recognize features of ENSO that are coherent, independent data series that represent each of these aspects have been used (Trenberth and Shea 1987). It is highly desirable to use the longest available data series for investigating interannual and decadal fluctuations, but data from the pre-1950 period are often missing or sparse (Wright 1989; Wang and Wang 1996). While this is the reason why many studies of ENSO concentrate on patterns from the 1950s onward, this recent period may not be representative of long-term ENSO dynamics (Trenberth and Shea 1987). The datasets used here have been filled and corrected carefully. The reliability of the pre-1950 data will be tested by comparing features of the series during that period.

The Southern Oscillation is the atmospheric component of ENSO. It is usually measured by the SOI, which is defined as the normalized monthly mean Tahiti minus Darwin sea level pressure anomalies. This form of the SOI is the most effective in capturing the out-of-phase relationship between pressure anomalies in the eastern and western Pacific (Trenberth 1984) and is the

most commonly used single-valued index of the Southern Oscillation. There is some concern over the quality of the early Tahiti data, which affects the quality of the SOI. Trenberth and Shea (1987) and Wang and Wang (1996) avoided the problem by only using the Darwin pressure anomalies. Since this may reduce the amount of ENSO-related variability being analyzed, we chose to use the SOI series of Allan (1993) in which the early part of the Tahiti record was corrected using data from neighboring stations. This SOI series extends from January 1876 to April 1995.

The oceanic component of ENSO can be measured by sea surface temperature (SST) anomalies in the equatorial Pacific Ocean. The data used here are anomalies of the gridded GISST 2.3a monthly SST set developed at the U.K. Meteorological Office (Rayner et al 1997). Average SST anomalies in the Niño3 region (5°N–5°S, 150°–90°W), where the ENSO signal is particularly high (Barnston 1995), are used (SST series hereafter). This series extends from January 1871 to December 1994. The pre-1982 SST data has been reconstructed using spatial EOFs. These EOFs were calculated using 45 yr of modern SST data and fitted in a least squares sense to the data available in each month from 1871 to 1981. No other climatic variables were used as input to this analysis. This regression technique is expected to improve the quality of SST data significantly, and certainly provides more interannual variance in the Niño3 region before 1950.

Another commonly used ENSO measure is the index of central equatorial Pacific rainfall (RAIN) constructed by Wright (1984, 1989). The record spans the period from October 1893 to August 1983. This index measures the tropical Pacific rainfall variations associated with the eastward movement of the region of tropical convection to the date line during a warm ENSO episode. A small number of missing data had been replaced by regression with sea surface temperature data (Wright 1989).

To increase the signal-to-noise ratio in the data, we used 3-month seasonal averages (December–February, March–May, etc.) of the original monthly data series. Such averages increase the signal-to-noise ratio of the data without affecting the interannual fluctuations in which we are interested (Trenberth 1984; Chu and Katz 1985). It should also be noted that since each of the three series represents anomalies from long-term mean monthly values, they contain almost no mean annual cycle.

The data series we have used are shown in Fig. 1. They are highly correlated for the periods of overlap between the series (see Table 1). Elliott and Angell (1988) and Trenberth and Shea (1987) have found that the correlation between the series has changed with time. Figure 2a shows a “moving correlation” between the SOI series and the SST and RAIN series, for 21-yr (84-season) overlapping segments of the data. It shows the same pattern as Elliott and Angell (1988) have found, namely that there is a drop in the correlation between each pair of series in the period 1920–50. The




FIG. 1. Time series of seasonal ENSO indices used in this study: Southern Oscillation index (SOI), Niño3 sea surface temperature (SST), and equatorial Pacific rainfall index (RAIN). Before plotting, the mean of each series was removed and the resulting anomalies were divided by the standard deviation of the series.

correlation is higher both before and after this period. A similar drop in the strength of the relationship occurs between other ENSO variables and teleconnections (Allan et al. 1995). It therefore appears that the weak relationship during 1920–50 represents a system-wide change, rather than due to weakening of some teleconnection mechanisms or poor data quality during this period (Elliott and Angell 1988).

The drop in the correlation between the indices coincides with a drop in the variance of each series (Elliott and Angell 1988; Trenberth and Shea 1987). Figure 2b shows the “moving variance” of 21-yr segments of the seasonal series. The plots show that the pre-1920 and post-1950 periods have a similar and large variance compared to the 1920–1950 period. This pattern has

TABLE 1. Correlation between the data series used in this study, for their periods of overlap. All values of the correlation coefficient are significant at 5% level.

	SOI	SST
SST	-0.68	—
RAIN	-0.66	0.80

also been found in several other ENSO series (Gu and Philander 1995).

The similarity of the changes in correlations and variance between the series prior to 1950 gives greater confidence in using these data in the early period, when their quality is more questionable. The similarity is greatest between the SOI and SST series, but is not as strong for the RAIN series.

3. Methods of time–frequency analysis

Time–frequency analysis, the spectral analysis of nonstationary signals whose frequency spectrum changes with time, is not trivial. Difficulties are caused by the properties of the time series becoming local rather than global in time. A truly local spectrum of a discrete nonstationary series exists only for the duration of one data point, but the estimation of a spectrum with one data point is virtually impossible. The usual method of dealing with this problem is to approximate the local behavior of the series at a point by the behavior of a short interval, or window, of data centered on this point. The window length is chosen

as a compromise between a better approximation of the local behavior (smaller windows) and better frequency resolution (larger windows with more data points). Consequently, there are no foolproof methods of time–frequency analysis.

We chose to analyze the ENSO data using three methods: windowed Fourier transform, wavelet transforms, and windowed Prony’s method. Each method has its own strengths and weaknesses. The comparison of their results can provide a more complete picture of ENSO’s time–frequency behavior.

Before analysis, each series of ENSO data had its long-term mean removed and was normalized by its standard deviation. This makes the total variance of each series equal to one and allows the spectra of the series to be compared directly.

a. Windowed Fourier transforms

The Fourier transform is a robust, well-understood, and often-used technique for spectral estimation. For time–frequency analysis the data were divided into equal-length (overlapping) segments (windows). The Fourier transform of each window was computed as an approximation to the spectrum of the point at the center of the window. This analysis technique is known as windowed Fourier transform, short-time Fourier transform, or the Gabor transform (Chan 1995). From here on it will be referred to as the Fourier method.

The Fourier method was carried out by performing a discrete Fourier transform on 21-yr (84-season) windows of the data. To prepare each window for spectral analysis, the mean of the window was subtracted, and the initial and final 5% of the data in the window was weighted with a cosine bell. The distance between the centers of neighbouring windows was 1 yr.

The problem with the Fourier transform stems from the fact that it only calculates the power of frequencies f_k , which belong to the discrete harmonic set

$$f_k \equiv \frac{k}{N\Delta t}, \quad k = 0, \dots, \frac{N}{2}, \quad (1)$$

where Δt is the sampling interval and N is the number of data points used in the analysis. This means that the resolution (or the number of power estimates) of the spectrum from a Fourier transform depends on the number of data points in the series. So while the Fourier transform works very well when performed on long data series, it may be of limited use for short data series, such as short data windows from a Fourier method analysis. The problem of the resolution is particularly obvious in the low-frequency part of the spectrum. Figure 3a illustrates how this problem affects the interannual part of a Fourier method spectrum (Fourier spectrum hereafter) of 21-yr windows of seasonal data. As can be seen in this figure, there are about five power estimates for periods in the range of 2–4 yr, and about the same number of estimations for periods in the range 4–




FIG. 2. Moving window analysis of overlapping 21-yr windows of ENSO series. The series were divided by their standard deviation for the period 1950–89 to emphasize differences in the early period. (a) Correlation (absolute value) of the SST and RAIN series with the SOI series. (b) Variance of the data segments. The value on the x axis indicates the year at the center of each segment.

21 yr. Clearly, the frequency resolution of ENSO-related periods is poor.

b. Wavelet analysis

Wavelet analysis is becoming increasingly common in meteorological studies. It is a method designed to optimize both time and frequency resolution by choosing the best window width for a particular frequency band. For this reason it is considered to be superior to the Fourier method and it is finding applications in many fields. The reader is referred to reviews by Farge (1992), Meyers et al. (1993), Lau and Weng (1995), and Torrence and Compo (1998) for detailed information on wavelets and their relationship to the Fourier method.

Briefly, the wavelet transform is defined by the convolution, or inner product, of the data signal $x(t)$ and a set of functions $\psi_{a,b}(t)$ called wavelets. Wavelets are

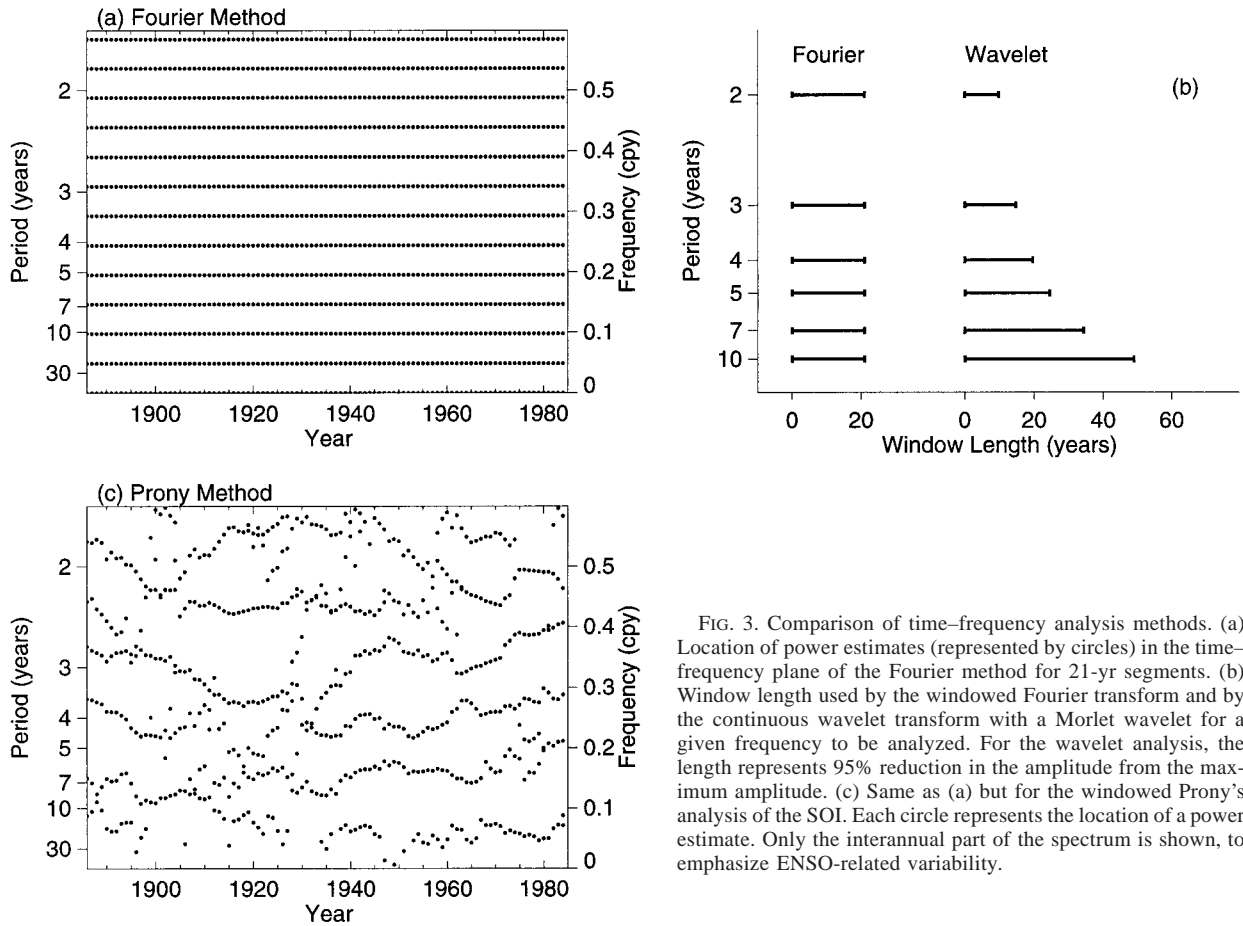


FIG. 3. Comparison of time-frequency analysis methods. (a) Location of power estimates (represented by circles) in the time-frequency plane of the Fourier method for 21-yr segments. (b) Window length used by the windowed Fourier transform and by the continuous wavelet transform with a Morlet wavelet for a given frequency to be analyzed. For the wavelet analysis, the length represents 95% reduction in the amplitude from the maximum amplitude. (c) Same as (a) but for the windowed Prony's analysis of the SOI. Each circle represents the location of a power estimate. Only the interannual part of the spectrum is shown, to emphasize ENSO-related variability.

derived from a “mother wavelet,” $\psi(t)$, by a translation by length b and a dilation by scale a :

$$\psi_{a,b}(t) = \frac{1}{a^{1/2}} \psi\left(\frac{t-b}{a}\right). \quad (2)$$

Wavelets have several special properties, one of which is compact support both in the time and the frequency spaces. A compactly supported function is effectively nonzero for a finite (and small) part of its domain. Hence wavelets can either be viewed as special windows in the time domain, or bandpass filters in the frequency domain (Chan 1995). The scale parameter a controls the width of the windows and the location of the bandpass filter and, hence, connects the window size to the resolution of particular frequencies. The parameter b controls the location of the window in the time domain.

The wavelet transform is defined by

$$W(a, b) = \frac{1}{a^{1/2}} \int \psi\left(\frac{t-b}{a}\right) x(t) dt, \quad (3)$$

where $W(a, b)$ is the wavelet coefficient at scale a and translation b . In the “continuous” wavelet transform (in contrast with the discrete wavelet transform) a can take

any positive value; its range of values is chosen according to the range of frequencies of interest in the data. The result is a clear graphical representation of how the spectrum changes with time (Farge 1992). Problems with this approach will be discussed in section 4. The values of b are limited to integers because of the discrete nature of the data.

The continuous wavelet transform (wavelet transform hereafter) can be efficiently calculated in Fourier space (Meyers 1993; Weng and Lau 1994) by using the following relationship:

$$W(a, b) = \frac{a^{1/2}}{2\pi} \int e^{-i\omega b} \hat{\psi}^*(a\omega) \hat{x}(\omega) d\omega. \quad (4)$$

The hat above a variable indicates the Fourier transform of that variable, that is,

$$\hat{\psi}(\omega) = \frac{1}{2\pi} \int \psi(t) e^{i\omega t} dt, \quad (5)$$

and the asterisk denotes the complex conjugate. Given a value of a , the product $\exp^{-i\omega b} \hat{\psi}^*(a\omega) \hat{x}(\omega)$ is calculated for each ω , and the result is inverse transformed using (4) to give $W(a, b)$ for each $b = 1, \dots, N$.

To minimize the distortion of the spectra due to end

effects, the series were padded on both sides with tails tapering to zero prior to being analyzed. The regions corresponding to the tails were discarded after the wavelet transform had been completed (Meyers et al. 1993).

A common way of displaying the results of the transform is to plot the amplitude of the wavelet coefficients $|W(a, b)|$ (Chan 1995; Farge 1992; Meyers et al. 1993). This method, however, is not directly comparable to the Fourier spectrum. For direct comparison of the spectra, we used instead the amplitude squared spectrum $|W(a, b)|^2$ (wavelet spectrum hereafter).

A mother wavelet that is commonly used in meteorology is the Morlet wavelet (Meyers et al. 1993; Gu and Philander 1995; Lau and Weng 1995; Wang and Wang 1996), which is defined by a complex sinusoid modulated by a Gaussian:

$$\psi(t) = e^{i\omega_0 t} e^{-(t^2/2)}. \quad (6)$$

In the Fourier domain the Morlet wavelet has the form

$$\hat{\psi}(\omega) = \frac{1}{2\pi^{1/2}} e^{-[(\omega - \omega_0)/2]^2}, \quad (7)$$

with the constant $\omega_0 = 6$ for the Morlet wavelet to ensure invertability (Farge 1992). Since the wavelet transform computes the similarity between the wavelets and the signal, the choice of the mother wavelet is important. Gu and Philander (1995) suggest that the ENSO signals may be represented by sinusoids, supporting the choice of the Morlet wavelet.

Figure 3b compares the length of the windows for the analysis of different frequencies between the Fourier method and the wavelet transform with the Morlet wavelet. It can clearly be seen that in the wavelet analysis the length of the window changes considerably between the high- and low-frequency parts of the spectrum, and uses only a similarly sized window to the Fourier method used here for periodicities in the 4–7-yr range.

c. Windowed Prony's method

Another way of overcoming the problem of poor frequency resolution from which the Fourier method suffers is to use a spectral analysis method that is not restricted to calculating the power of a particular set of frequencies. The Prony method of spectral analysis (e.g., Marple 1987; Barone et al. 1989), for example, assumes that the data can be modeled as the linear sum of exponentially decaying sinusoids:

$$\hat{x}_k = \sum_{n=1}^q A_n \exp(\alpha_n k \Delta t) \cos(2\pi f_n \Delta t + \phi_n), \quad k = 0, \dots, N - 1, \quad (8)$$

where \hat{x}_k is the model approximation of order q to the data series x_k , ($k = 0, \dots, N - 1$). The parameters A_n , α_n , f_n , and ϕ_n are the amplitudes, decay factors, frequencies, and phases of the model, respectively. Unlike

Fourier transforms, the choices of the frequencies in the Prony method are not restricted a priori.

Briefly, the Prony method calculates the $4q$ model parameters, which minimize the squared error

$$\sum_{k=0}^{N-1} |x_k - \hat{x}_k|^2, \quad (9)$$

given a model order q . The solution of this nonlinear least squares problem reduces to fitting an autoregressive (AR) model of order $2q$ to the data, and the autoregressive coefficients are then used to calculate the frequencies and decay factors. The amplitudes and phases are then found through the solution of a linear least squares problem.

This method, applied to 21-yr segments of the data as was done with the Fourier method, from here on shall be referred to as the Prony method. We have fitted each window of the data with the maximum model order q possible, which is the maximum q such that $4q \leq N$ (i.e., $q = 21$ for seasonal data). Figure 3c shows the same part of the time–frequency plane as Fig. 3a, but with the results of a Prony frequency analysis. The circles directly above each year on the x axis are the q frequency components obtained from a 21-yr segment centered on that year. This demonstrates how the frequencies resolved are not restricted to a fixed discrete set of frequencies, as in the Fourier method.

If the true order for the data is less than the order q chosen, then some of the frequency components resolved by the method are the result of overfitting the data and are of little interest. To assess the importance of each frequency component f_n to the data, the damped sinusoid component with coefficients $\{A_n, \alpha_n, f_n, \phi_n\}$ was reconstructed using the equation

$$\hat{x}_k(n) = A_n \exp(\alpha_n k \Delta t) \cos(2\pi f_n \Delta t + \phi_n), \quad k = 0, \dots, N - 1. \quad (10)$$

Then, $\hat{x}_k(n)$ was cross correlated with the original data x_k and the correlation coefficient squared r_n^2 (fraction of variance explained by the n th frequency component) was used as a measure of the strength of this component. A frequency component that features strongly in the data will have a higher value of r_n^2 than one that does not. To facilitate plotting the Prony spectrum as a contour plot, the discrete deltalike pulses of the Prony spectrum were converted to narrow triangular functions in the frequency plane with a peak at height $r_n^2 \times$ (variance of window) and base at zero.

4. Results

a. Comparison of time–frequency methods

The Fourier, wavelet, and Prony time–frequency spectra of the SOI are shown in Fig. 4. Although they appear qualitatively different, a closer inspection shows that the spectra have a similar power distribution, par-

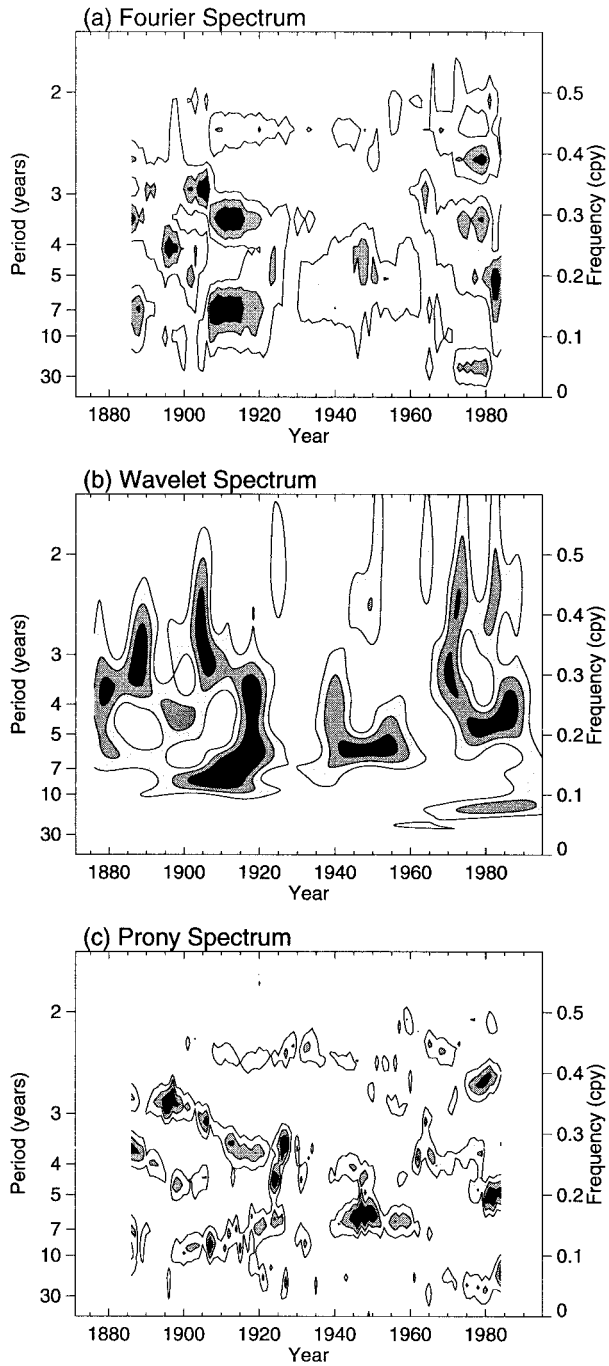


FIG. 4. Time–frequency spectrum of the SOI using (a) the Fourier method, (b) continuous wavelet analysis, and (c) windowed Prony’s method. In (a) and (c) the value on the x axis represents the year at the center of a 21-yr window. For each spectrum, contour lines represent 20%, 40%, and 60% of the maximum power of that spectrum. The black shading therefore represents regions where the power is greater than 60% of the maximum.




FIG. 5. Time-integrated and standardized spectra of the wavelet (solid line) spectrum and the Fourier spectrum (broken line) of the SOI.

ticularly in the location of strong power. The variability of the SOI is contained in a 2–10-yr periodicity band and there are large decadal changes in the power distribution within this band. There is little evidence for any fixed frequency or mode that is persistent throughout the period of analysis. There are, however, some important differences between the methods that need to be discussed before the behavior of the SOI and the other series can be described more carefully.

The Fourier method has an inherent tendency to overestimate power at high frequencies at the expense of low frequencies, while the wavelet transform improves this bias (Lau and Weng 1995). This is illustrated in Fig. 5 by comparing the time-averaged Fourier and wavelet spectra, created from the time–frequency spectra. Apart from differences in the frequency resolution (which is why the peaks of the Fourier spectrum and the wavelet spectrum are slightly misaligned), the spectra are very similar. Two notable differences exist: the QB peak of the Fourier spectrum at approximately 0.43 cpy is barely noticeable in the wavelet spectrum, and the peak in the wavelet spectrum at approximately 0.08 cpy is barely noticeable in the Fourier spectrum. Since the wavelet spectrum has almost no QB peak, it is possible that Fourier analysis in general may produce spurious QB peaks in ENSO series, particularly if short series are analyzed.

The fixed window length of the Fourier and Prony methods can also produce misleading results. The Fourier and the Prony spectra in Figs. 4a and 4c, respectively, have what may be called a QB band, represented by a persistent peak with a period of about 2.4 yr. This persistence, however, is probably an artifact of poor time localization because in the wavelet spectrum these 2-yr cycles occur as pulses (e.g., at approximately 1925, 1945, 1964, 1972, and 1982). The 21-yr windows used with the Fourier and Prony methods are much larger than the period of the oscillation so a 2-yr pulse remains

within the analyzing window for a long time. The wavelet transform uses a window approximately 10 yr long, as is shown in Fig. 3b, and hence has superior time localization. A plausible explanation for the pulses with a period of approximately 2 yr is the tendency for an El Niño episode to be preceded or followed by a La Niña episode. Each of these episodes lasts for about 1 yr, creating overall a cycle with a period of about 2 yr (Barnett 1991).

The wavelet and Prony spectra have improved frequency resolution at low frequencies compared to the Fourier spectrum. For example, the Fourier spectrum shows an apparently constant 7-yr periodicity from 1905 to 1920, because it does not resolve any periods between 7 and 10 yr. Both the wavelet and the Prony methods indicate that over this period, the period of oscillation has gradually changed from about 10 yr to about 7 yr.

In the wavelet spectrum in Fig. 4b the spectral peaks at high frequencies tend to be “stretched” in the frequency plane. This is a consequence of using continuous wavelet analysis. If the covariance of a signal with a sinusoid of a particular frequency is high, then the covariance of the signal with a sinusoid with a neighboring frequency will also be high. Another way of looking at this problem is that the wavelet spectrum is “overcomplete.” Because we can analyze the data with as many values of the scale parameter a as we choose, the wavelet spectrum contains many more data values (but no more information) than the original data. The extra data values create broader peaks and smeared power, and reduces the certainty in the exact position of a particular spectral peak (Farge 1992). We overcame this problem by comparing the wavelet spectrum to those produced by the Fourier and Prony methods, both of which produce only as much spectral information as that contained in the data.

Wang and Wang (1996) analyzed Darwin sea level pressure anomalies using continuous wavelet analysis. Their resulting spectrum is very similar to the spectrum of the SOI in Fig. 4b (Fig. 8a in Wang and Wang 1996, 1595). To overcome the problems inherent in the wavelet analysis, they also used an “improved” method called waveform analysis. The waveform spectrum was also clearly nonstationary. However, it produced features that appear to be constant frequency modes that lasted for a long period of time, like a strong 3-yr oscillation during 1890–1910 (Fig. 8b in Wang and Wang 1996, 1595). Evidence from the spectra of the SOI in this study does not support these conclusions. The strong oscillations are usually short lived and change their frequency continuously. Within the limitations of each of the time–frequency methods used here, they all agree on this point. The problems with the waveform method used by Wang and Wang (1996) may be that it has low-frequency resolution (even lower than the Fourier method here) and that, as Wang and Wang (1996) noted, it has poor amplitude estimation capabilities.

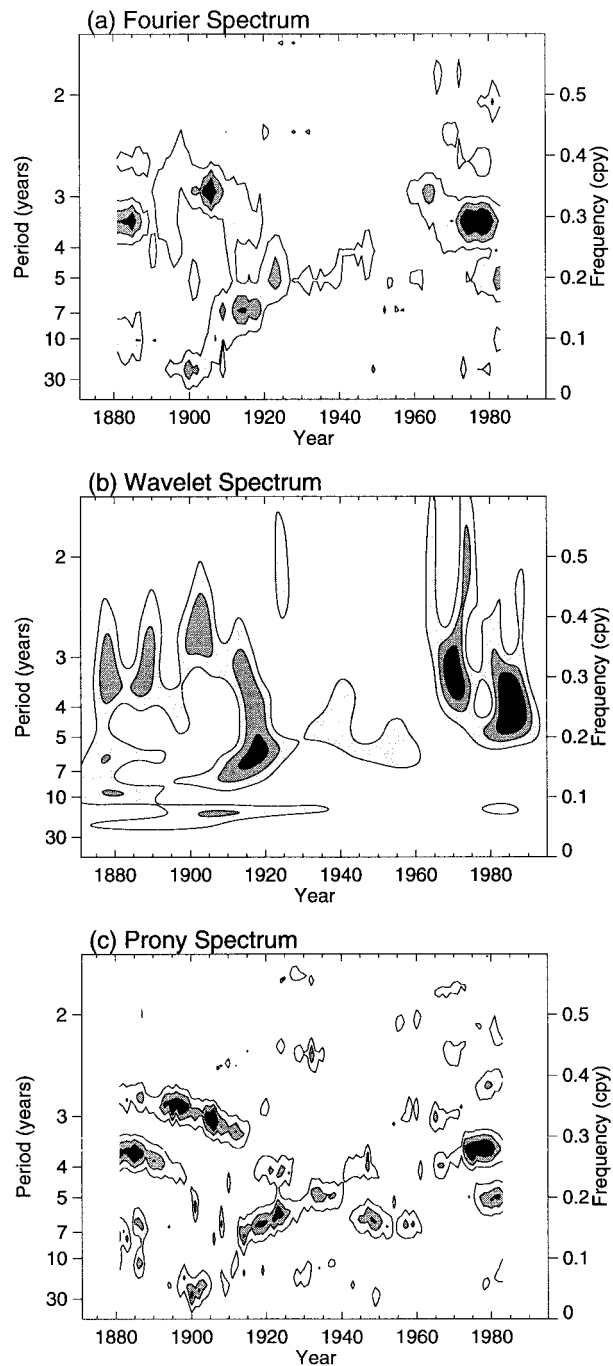


FIG. 6. Same as Fig. 4 but for the SST series.

b. Comparison of spectral patterns of ENSO series

The time–frequency spectra of the SST and the RAIN series are given in Figs. 6 and 7, respectively. The qualitative differences between the analysis methods, noted above for the SOI, are also apparent for the SST and RAIN series and will not be elaborated on further.

By inspection, the spectra of the SOI, SST, and RAIN series have very similar patterns. An objective measure

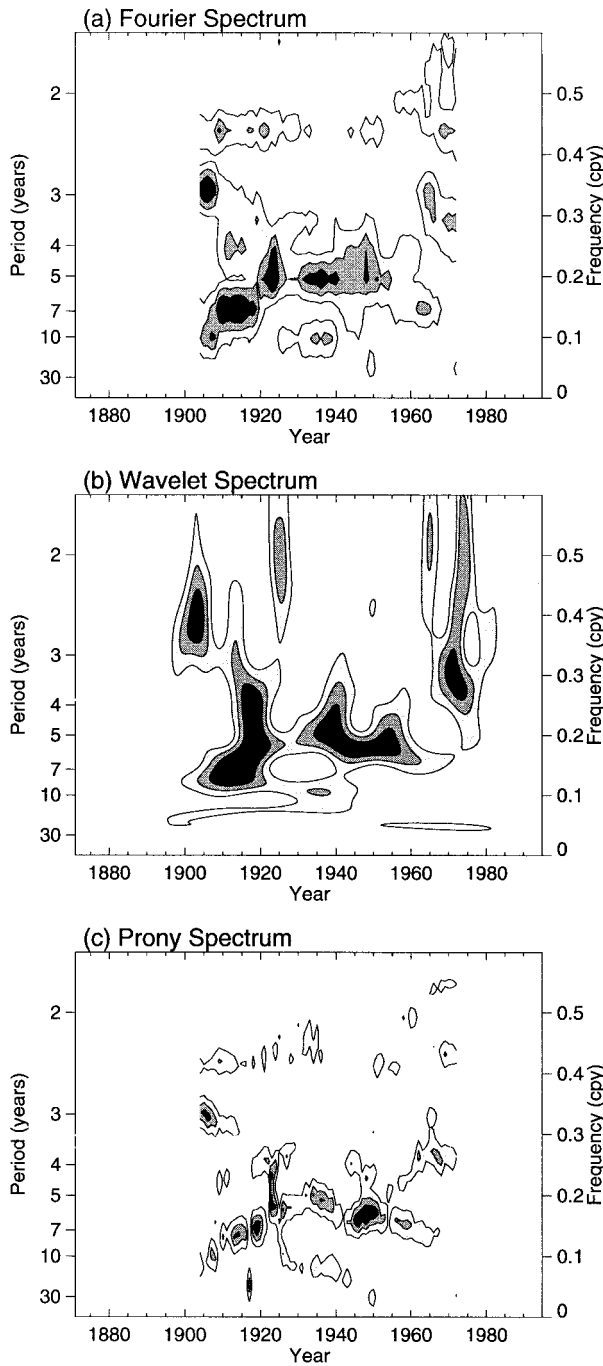


FIG. 7. Same as Fig. 4 but for the RAIN series.

TABLE 2. Pattern congruences between the wavelet amplitude squared spectra of the SOI, SST, and RAIN series, for their whole periods of overlap.

	SOI	SST
SST	0.86	—
RAIN	0.90	0.89

TABLE 3. Same as for Table 2 but for the pre-1950 (underlined) and post-1950 periods of overlap.

	SOI	SST
SST	<u>0.85/0.88</u>	—
RAIN	<u>0.89/0.90</u>	<u>0.90/0.89</u>

of the similarity between the spectra can be obtained by calculating the pattern congruence [uncentered pattern correlation, Harman (1976)] between pairs of spectra; here we compare the wavelet spectra. For two time–frequency spectra X_{ij} and Y_{ij} for $i = 1, \dots, N, j = 1, \dots, M$ the pattern congruence is defined by

$$C = \frac{\sum_{i,j} X_{ij} Y_{ij}}{\left(\sum_{i,j} X_{ij}^2 \sum_{i,j} Y_{ij}^2 \right)^{1/2}} \quad (11)$$

A value of C close to one means that the spectral patterns are very similar, while a value of C close to zero means that the spectra are not similar.

The values of the congruence between the wavelet spectra of the SOI, SST, and RAIN, for their entire period of overlap, are given in Table 2. These values are close to one and indicate that the data series have similar spectral structures. As yet we cannot assess how statistically significant these congruence values are. In section 5 we will develop a test for significance using stochastic models of the SOI.

It was mentioned in section 2 that the pre-1950 data is often treated with suspicion because its quality may be low. Table 3 compares the values of the congruence for the pre-1950 and post-1950 periods of overlap between the data series. The results show that the spectra of the series are about as similar in the pre-1950 period as they are in the post-1950 period. As the data series were mostly derived independently (see section 2), the agreement between their spectral structures and temporal structures means that the series capture the essence of the same phenomenon—ENSO. The agreement increases our confidence in the reliability of the data in the pre-1950 period, and in the validity of the spectral patterns that have been calculated.

The distribution of the spectral power of frequencies in time is episodic, rather than continuous, as Figs. 4, 6, and 7 show. From the late 1800s to 1925, the dominant frequencies occur in the 3–10-yr range. This is not a single coherent mode, and it is not so clearly separated from higher frequencies in the 2–3-yr range as is often thought (e.g., Barnett 1991; Wang and Wang 1996). During the period 1930–60, the presence of higher frequencies is very weak, and the spectrum is dominated by a low 4–7-yr periodicity. This period approximately corresponds to the period of reduced variance and correlation between ENSO series (section 2b; Allan 1993; Trenberth and Shea 1987; Cole et al. 1993). In the 1970s the period of the fluctuations decreases again

to 2–5 yr. These results, together with the observed changes in the variance of ENSO, support the assertion that the recent activity of ENSO is more similar to the late 1800s to early 1900s than to the period in the middle of the twentieth century (Allan 1993).

The SST series (GISST 2.3a) provides a marked improvement in quality over the frequently used SST data from the Comprehensive Ocean Atmosphere Data Set (COADS). Wang and Wang (1996) analyzed central and eastern equatorial Pacific COADS SST data using wavelet analysis. They found that the spectrum from this data compared poorly with the Darwin sea level pressure anomaly series in the pre-1950 period. As mentioned earlier, the Darwin pressure series produced similar spectral patterns to the series used in this investigation and we infer that the COADS SST data is of lower quality in the pre-1950 period.

5. Statistical testing with stochastic simulations of the SOI

Power spectra and other properties of time series are usually tested statistically against white noise or red noise processes. A more rigorous test can be provided by using a stochastic process, like an autoregressive-moving average (ARMA) model, which imitates properties of the series (Chu and Katz 1985; Trenberth and Hoar 1996; Torrence and Compo 1998). In this section, an ARMA model will first be used to assess the significance of the congruence test statistic. It will then be used to assess the significance of a temporal and a spectral ENSO “trademark” (variance and wavelet spectrum, respectively). In addition, we can generate simulated series by randomizing the phase of the individual spectral components of the observed SOI and then recombining them to generate new series. While the new series have the same power spectrum as the SOI, they may have a different temporal structure, and hence a different time–frequency spectrum, because of the random phase relationships. This procedure has been used by Yamada and Ohkitani (1991) and Ivanov et al. (1996) to test for nonrandom phase structures in time series, which may indicate a singularity in the time series.

Chu and Katz (1985) found that the best ARMA model to fit the seasonal series of the SOI from 1935 to 1983 is the AR(3) model specified by

$$X_t = 0.6885X_{t-1} + 0.2460X_{t-2} - 0.3497X_{t-3} + a_t \tag{12}$$

Here, X_t is the value of the process at time t , and a_t constitutes a white noise process with variance $\sigma_a^2 = 1.505$ [Eq. (12) will be referred to from here on simply as the AR(3) model]. Trenberth and Hoar (1996) found an ARMA(3,1) model for Darwin sea level pressure anomalies for the period 1882–1981. While they have used a longer record than Chu and Katz (1985) to calculate their model, the Chu and Katz (1985) model was used here because it uses the SOI directly, and is not




FIG. 8. Theoretical power spectrum of the AR(3) model of the SOI (light line) together with the spectrum of the real SOI series for the period 1935–83 (solid heavy line) and for the whole record 1876–1995 (broken heavy line).

based on data that was prefiltered. In general, both models behave similarly.

The deterministic component of the AR(3) model, obtained by removing the white noise forcing, is a decaying oscillation with a period of about 3 yr. Stochastic forcing disturbs the decaying oscillator so that the spectral peak is broadened and the process continues. The AR model was designed to imitate the autocorrelation structure of the SOI, and as a consequence it imitates the spectrum of the SOI. The theoretical spectrum of the AR(3) model, calculated using model coefficients (Jenkins and Watts 1968), is plotted in Fig. 8. Also plotted in Fig. 8 are spectra of the SOI that demonstrate that the spectrum of the AR(3) model closely resembles the general shape of the spectrum of the SOI for periods up to about 10 yr. However, for periods longer than 10 yr the model has more power than the observed SOI.

A realization of the AR(3) model is plotted in Fig. 9a. It was created by seeding (12) with random initial conditions and allowing the process to run for 1478 iterations. The first 1000 iterations were discarded to eliminate the effects of the initial conditions (Trenberth and Hoar 1996). The long-term mean of the resulting 478-season long time series (length of the real SOI series) was removed and the series was normalized by its long-term standard deviation. This procedure cast the series in a form directly comparable to the standardized ENSO series described in section 3. Many realizations can be created by repeating this process with different random initial conditions. The wavelet spectrum of the realization in Fig. 9a is plotted in Fig. 9b. The qualitative resemblance between the spectral patterns produced by the synthetic SOI and those produced by the real SOI in Fig. 4b is striking. Torrence and Compo (1998) have shown that an AR(1) red noise process can also simulate much of the time–frequency variability shown in the wavelet power spectrum. Since the AR(3) model is sto-

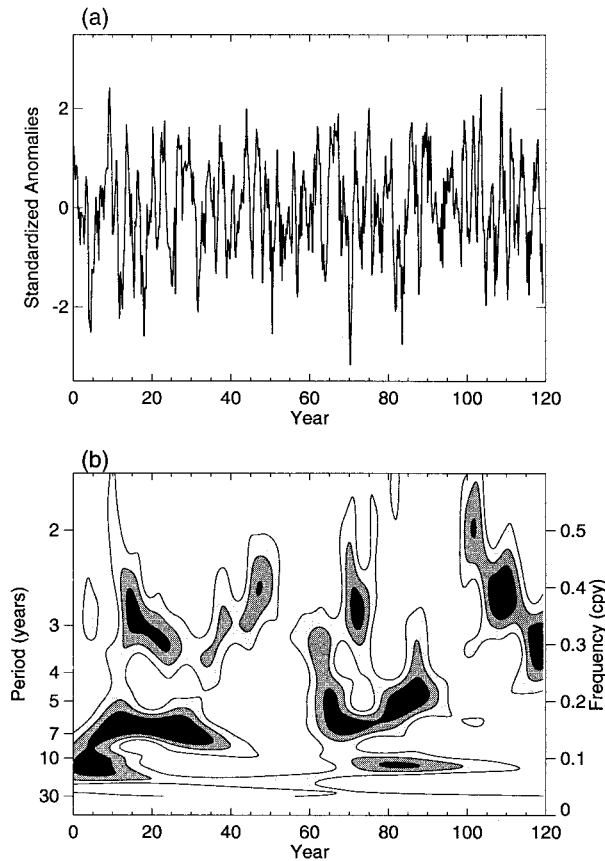


FIG. 9. A realization of the AR(3) model with random initial conditions: (a) time series plot and (b) wavelet amplitude squared spectrum of the realization, contoured as in Fig. 4.

chastic, the peaks in its wavelet spectrum have no deterministic origin and are called pseudofrequencies (Chu and Katz 1985). They occur when the model behaves approximately periodically for a short time.

a. Pattern congruences

In section 3b the similarity between the wavelet spectra of the SOI, SST, and RAIN series was measured using pattern congruences. The congruence values in Table 2 and Table 3 are high, but may not be significant if the congruence between random series is also of comparable magnitude. To test the significance, we calculated the congruence between the wavelet spectra of 100 pairs of series randomly selected from a pool of 200 realizations of the AR(3) model. The length of the series equaled the shortest length of overlap period between the series tested, which is 359 seasons. Here, 95% of the pairs had a congruence of less than 0.63. A similar test was carried out with the phase-randomized series, by calculating the congruence between the observed SOI series and 100 phase-randomized series. In this test 95% of the the congruence values were less than 0.67. Both these tests indicate that the time–frequency spectra

TABLE 4. Maximum and minimum variance (in standard deviation units) of 84-season-long data segments from a 478-season-long data series. The values in the square brackets indicate the 95% confidence intervals calculated from 1000 realizations of simulated SOI series. All the series had been standardized as described in section 3 prior to analysis.

Variance	Min	Max
Observed SOI	0.48	1.50
AR(3) model	[0.42, 0.74]	[1.23, 1.82]
Phase-randomized SOI	[0.44, 0.76]	[1.22, 1.76]

of the observed ENSO series are all significantly similar to each other at the 5% level (see Table 2).

b. Range of variance

The variance of the SOI has changed considerably during the past century (see Fig. 2b and section 2b). These changes have been large enough to be attributed to decadal variations in the background state of the atmospheric circulation (Elliott and Angell 1988; Allan 1993; Allan et al. 1995). In Table 4 the minimum and maximum values of the variances of 21-yr windows of the SOI are compared to the values produced by AR(3) simulations and phase-randomized SOI series. One thousand realizations of each of the simulated series were used to produce two-tailed 95% confidence intervals for the maximum and minimum variances. The values of the minimum and maximum variance of the observed SOI fall within the confidence intervals of both the AR(3) series and the phase-randomized series. This means, first, that phase-randomization has not altered the temporal structure of the SOI significantly. Second, neither the maximum nor the minimum variance of the real SOI is significantly different from what would be expected if it was a realization of the AR(3) model. The AR(3) model is therefore able to produce decadal changes in variance that are similar to those produced by the SOI.

c. Distribution of wavelet coefficients

The average spectrum of the SOI series has a broad spectral peak (Fig. 8), but the distribution of power over time is more sharply peaked in frequency and is not constant [Figs. 4, 6, and 7; Wang and Wang (1996)]. The persistence and distribution of power in the AR(3) spectrum (Figs. 8 and 9) are very reminiscent of those observed for the SOI and the other ENSO series. To test whether there are significant differences in structure between the real and synthetic SOI spectra, the probability distribution of wavelet coefficient amplitudes was calculated. This technique has been used before with discrete wavelet analysis (Yamada and Ohkitani 1991; Ivanov et al. 1996) to distinguish between wavelet spectra with different structures. For example, Yamada and Ohkitani (1991) found that while the distribution of

wavelet coefficients of a phase-randomized turbulence time series was Gaussian, the turbulence series itself had an exponential distribution. The distribution of the wavelet coefficients of phase-randomized series is used as a null hypothesis for a completely random, structureless series.

The wavelet spectra were analyzed as follows. Each wavelet spectrum was divided into four frequency bands: 0.5–0.4, 0.4–0.3, 0.3–0.2, and 0.2–0.1 cpy (2–2.5, 2.5–3.3, 3.3–5, and 5–10 yr, respectively). Four histograms of wavelet coefficients $|W_{a,b}|^2$ were obtained by sorting the coefficients within each band according to their magnitude. This analysis was repeated 1000 times with 478-season-long AR(3) realizations and with 478-season-long phase-randomized series. The multiple analyses were used to estimate the mean and standard deviation of histogram heights for each frequency band of the AR(3) model and the phase-randomized SOI. Finally, the histogram values were converted to probability densities.

Figure 10 shows the probability distribution of the squared amplitude of wavelet coefficients $|W_{a,b}|^2$. The distribution of the real SOI series is also plotted. The distribution of the real series resembles the distribution of the AR(3) model in the 2–2.5-, 2.5–3.3-, and 5–10-yr bands (Figs. 10a, 10b, and 10d), and is well within the confidence limits. In the 3.3–5-yr frequency band (Fig. 10c) the distribution is more erratic and one excursion outside the confidence limits occurs. In the 10–50-yr band (not shown) the model has appreciably more large amplitude coefficients than the observed SOI because the model overestimates the power in this frequency band.

The distributions of the AR(3) model in Fig. 10 are decaying exponentials, which means that the distribution of $|W_{a,b}|$ (unsquared coefficients) was Gaussian. This type of distribution is expected from a regular stochastic process like the AR(3) model (Yamada and Ohkitani 1991). It appears that, particularly in the high-frequency range (2.0–3.3 yr) the SOI also behaves like a stochastic system. The resemblance to a decaying exponential is not as strong for the 5–10-yr band, though still significantly similar. Something different appears to be occurring in the 3.3–5.0-yr band. Power estimates with an amplitude of about 0.3 occur more than would be expected had the SOI behaved exactly like the AR(3) model. This could be caused, for example, by a dynamical mechanism operating at this timescale.

The distribution of the observed SOI does not fit the distribution of the phase-randomized SOI particularly well. While the phase-randomized distributions are also decaying exponentials, they have somewhat different distributions to those of the AR(3) and the SOI. The difference may be attributed to the autocorrelation structure of the AR(3) model and the observed SOI series.

In general, the simple AR(3) model described in (12) is capable of producing both temporal and spectral decadal variabilities that are very similar to those observed

in the real seasonal SOI data series over the past 120 yr. It does so without any change of background state (like a change in the variance of forcing). This may indicate that the SOI, and hence other ENSO indicators, are driven by a similar stochastic mechanism.

6. Discussion and conclusions

A fundamental gap in our understanding of the ENSO phenomenon is the uncertainty about which underlying mechanisms, dynamical or stochastic, are primarily responsible for its observed variability (Chang et al. 1996). This knowledge is crucial for efficient and accurate modeling and prediction of ENSO. The lack of long-term high quality data has been one of the major hindrances for progress in this area. In this study we investigated how the temporal and spectral properties of observed ENSO series have changed over the period of observation. The seasonal time series used—SOI, Niño3 SST, and central equatorial Pacific rainfall index—are the longest available quality ENSO series.

The temporal structure of the series varied considerably over the past century. All the series had low variance in the period 1920–50, compared with their variance in the periods immediately beforehand and afterward. The period of low variance coincided with a period of reduced correlation between the series. These patterns in temporal structure have also been observed for other ENSO series, suggesting that these changes were a system-wide phenomenon (Elliott and Angell 1988).

The interannual variability of ENSO is contained in a broad 2–10-yr periodicity band. By using time–frequency analysis methods, which analyze changes in the spectrum of the ENSO series, it was found that this variability is not distributed evenly over time, but rather that it is concentrated at different frequencies at different times (Wang and Wang 1996; Cole et al. 1993). For example, the period from 1930 to 1960 is dominated by a 4–7-yr periodicity. In the earlier period (1880–1930) and the later period (1960–90) shorter periods of 2–5 yr also occur, strengthening suggestions that the current state of ENSO is similar in magnitude and frequency to the conditions late last century and early this century (Allan 1993). However, concentrated power does not appear to be fixed to either certain frequencies or to distinct frequency bands, as has been suggested previously (Barnett 1991; Wang and Wang 1996).

The temporal and spectral patterns described above are significantly similar between all the series, particularly between the SOI and SST series. It supports the assumption that these series represent aspects of the same phenomenon, and increases the confidence in the quality of the data series, even in the pre-1950 period. These results open the way for more long-term studies of ENSO and the use of the extended SOI and GISST 2.3a datasets for the calculation of ENSO-related climatologies. We believe that these applications are par-

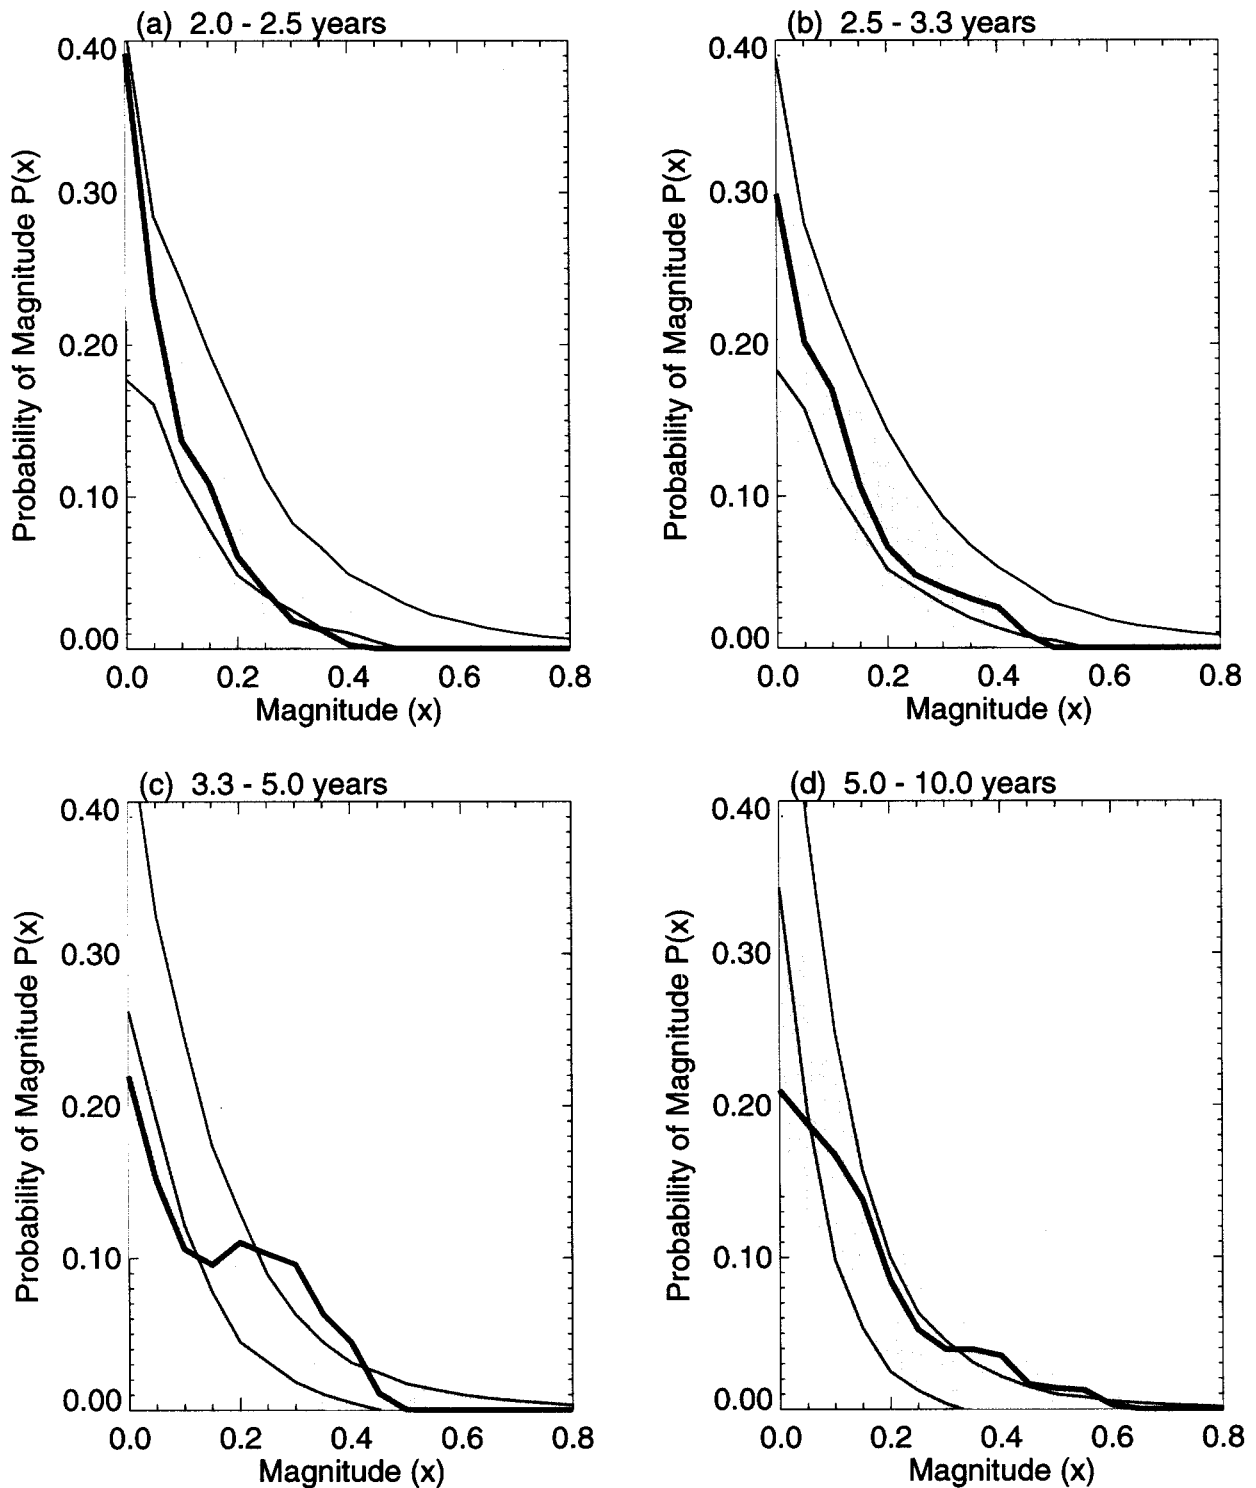


FIG. 10. The probability density function of wavelet coefficients in the (a) 2.0–2.5-yr band, (b) 2.5–3.3-yr band, (c) 3.3–5.0-yr band, and (d) 5.0–10.0-yr band. The shaded region corresponds to the mean \pm two standard deviations of 1000 realizations of the AR(3) model. The fine solid lines are the mean \pm two standard deviations of 1000 realizations of the phase-randomized SOI. The heavy solid lines are the distributions of the wavelet coefficients for the observed SOI from Fig. 4b.

ticularly important because, as can be seen in the spectra in Figs. 4, 6, and 7, the period from 1950 onward gives an incomplete view of the full complexity of ENSO (Trenberth and Shea 1987). It would be interesting to see whether the results of studies that examined only post-1950 ENSO data are reproducible when extended sets are used.

The wavelet coefficients of the observed SOI have a decaying exponential distribution over most of the frequency range. This indicates that the SOI behaves like a stochastic system (Yamada and Ohkitani 1991). It does not appear to be an entirely random process, as its distribution is (at least in the parameter values of the distribution) significantly different from the distribution created by phase-randomized SOI series. However, the distribution of the observed SOI is within the confidence bounds of a distribution created from simulated series of an AR(3) model that was fitted to the SOI by Chu and Katz (1985). It only deviated from this model in the 3.3–5.0 periodicity band. The deviation in this periodicity band may reflect some underlying dynamics of ENSO. The AR(3) series were also able to imitate the decadal changes in the variance that have been observed for the SOI. It appears, therefore, that the simple stochastic AR(3) process can imitate the key features of both the temporal and spectral interannual variability of the SOI, without any change to the background forcing.

We emphasize that the above conclusion does not mean that the ENSO system is an autoregressive process of order 3. The model used was empirical and was not based on knowledge of the physical system. The comparison between the model and ENSO provides evidence that the similarity could result from ENSO being mostly a well-behaved stochastic system. The analogy might be extended further if we remember that the AR(3) model [and the ARMA(3,1) model of Trenberth and Hoar (1996)] is a linear oscillator driven by stochastic forcing. The above results may indicate that the ENSO system belongs to this class of systems. This was suggested by Penland and Sardeshmukh (1995) who found, using inverse modeling of SST data, that ENSO was approximated well by a linear system forced by white noise. However, this discussion must be assessed in light of Rajagopalan et al. (1997) who cautioned that the short length of ENSO data available may increase the sensitivity of the results to how the analysis was carried out.

Support for the stochastic origin of interannual ENSO variability has been growing. Several recent studies (Kleeman and Moore 1997; Chang et al. 1996; Flügel and Chang 1996; Jin et al. 1996) have found that stochastic processes are the major source of ENSO variability. Further evidence is provided by the fact that the forecasting skill of dynamical models based on our physical knowledge of the ENSO system is no better than that of empirical models derived from data series (Barnston 1995). This suggests that our understanding of the physics of the system is far from perfect, or that

the predictability of the models may be inherently limited by stochastic forcing (Kleeman and Moore 1997; Flügel and Chang 1996). Here we presented evidence that the variability of observed ENSO data is not inconsistent with a stochastic origin.

Acknowledgments. We are grateful to Rob Allan and Phil Jones for the provision of data for this study. For helpful discussions and comments we thank Rob Allan, Richard Kleeman, Wasyl Drosdowsky, Bill Randel, Chris Folland, Zoltan Toth, and an anonymous reviewer.

This study was supported by the the Australian Government through the Cooperative Research Centres program.

REFERENCES

- Allan, R. J., 1993: Historical fluctuations in ENSO and teleconnection structure since 1879: Near-global patterns. *Quat. Australasia*, **11**, 17–27.
- , J. A. Lindesay, and C. J. C. Reason, 1995: Multidecadal variability in the climate system over the Indian Ocean region during the austral summer. *J. Climate*, **8**, 1853–1873.
- Barnett, T. P., 1991: The interaction of multiple time scales in the tropical climate system. *J. Climate*, **4**, 269–285.
- , M. Latif, N. Graham, and M. Flügel, 1995: On the frequency–wavenumber structure of the tropical ocean/atmosphere system. *Tellus*, **47A**, 998–1012.
- Barnston, A. G., 1995: Our improving capability in ENSO forecasting. *Weather*, **50**, 419–430.
- Barone, P., E. Massaro, and A. Polichetti, 1989: The segmented Prony method for the analysis of non-stationary time series. *Astron. Astrophys.*, **209**, 435–444.
- Brassington, G. B., 1997: The modal evolution of the Southern Oscillation. *J. Climate*, **10**, 1021–1034.
- Chan, Y. T., 1995: *Wavelet Basics*. Kluwer Academic, 134 pp.
- Chang, P., L. Ji, H. Li, and M. Flügel, 1996: Chaotic dynamics versus stochastic processes in El Niño–Southern Oscillation in coupled ocean–atmosphere models. *Physica D*, **98**, 301–320.
- Chu, P.-S., and R. W. Katz, 1985: Modeling and forecasting the Southern Oscillation: A time–domain approach. *Mon. Wea. Rev.*, **113**, 1876–1888.
- Cole, J. E., R. G. Fairbanks, and G. T. Shen, 1993: Recent variability in the Southern Oscillation: Isotopic results from Tarawa Atoll coral. *Science*, **260**, 1790–1793.
- Elliott, W. P., and J. K. Angell, 1988: Evidence for changes in Southern Oscillation relationships during the last 100 years. *J. Climate*, **1**, 729–737.
- Farge, M., 1992: Wavelet transforms and their applications to turbulence. *Annu. Rev. Fluid Mech.*, **24**, 395–457.
- Flügel, M., and P. Chang, 1996: Impact of dynamical and stochastic processes on the predictability of ENSO. *Geophys. Res. Lett.*, **23**, 2089–2092.
- Gu, D., and S. G. H. Philander, 1995: Secular changes of annual and interannual variability in the Tropics during the past century. *J. Climate*, **8**, 865–876.
- Harman, H. H., 1976: *Modern Factor Analysis*. University of Chicago Press, 487 pp.
- Ivanov, P. Ch., M. G. Rosenbaum, C.-K. Peng, J. Mietus, S. Havlin, H. E. Stanley, and A. L. Goldberger, 1996: Scaling behaviour of heartbeat intervals obtained by wavelet-based time-series analysis. *Nature*, **383**, 323–326.
- Jenkins, G. M., and D. G. Watts, 1968: *Spectral Analysis and Its Applications*. Holden-Day, 525 pp.
- Jin, F.-F., J. D. Neelin, and M. Ghil, 1996: El Niño/Southern Oscillation and the annual cycle: Subharmonic frequency-locking and aperiodicity. *Physica D*, **98**, 442–465.

- Kleeman, R., and R. Moore, 1997: A theory for the limitations of ENSO predictability due to stochastic atmospheric transients. *J. Atmos. Sci.*, **54**, 753–767.
- Lau, K.-M., and H.-Y. Weng, 1995: Climate signal detection using wavelet transform: How to make a time series sing. *Bull. Amer. Meteor. Soc.*, **76**, 2391–2402.
- Marple, S. L., 1987: *Digital Spectral Analysis with Applications*. Prentice-Hall, 492 pp.
- Meyers, S. D., B. G. Kelly, and J. J. O'Brien, 1993: An introduction to wavelet analysis in oceanography and meteorology: With application to the dispersion of Yanai waves. *Mon. Wea. Rev.*, **121**, 2858–2866.
- Penland, C., and P. D. Sardeshmukh, 1995: The optimal growth of tropical sea surface temperature anomalies. *J. Climate*, **8**, 1999–2024.
- Rajagopalan, B., U. Lall, and M. A. Cane, 1997: Anomalous ENSO occurrences: An alternate view. *J. Climate*, **10**, 2351–2357.
- 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.
- Rayner, N. A., E. B. Horton, D. E. Parker, C. K. Folland, and R. B. Hackett, 1997: Version 2.2 of the global sea-ice and sea surface temperature data set 1903–1994. CRTN 74. [Available from Hadley Centre for Climate Prediction and Research, London Rd., Bracknell, Berkshire RG12 2SY, United Kingdom.]
- Torrence, C., and G. P. Compo, 1998: A practical guide to wavelet analysis. *Bull. Amer. Meteor. Soc.*, **79**, 61–78.
- Trenberth, K. E., 1984: Signal versus noise in the Southern Oscillation. *Mon. Wea. Rev.*, **112**, 326–332.
- , and D. J. Shea, 1987: On the evolution of the Southern Oscillation. *Mon. Wea. Rev.*, **115**, 3078–3096.
- , and T. J. Hoar, 1996: The 1990–1995 El Niño–Southern Oscillation event: Longest on record. *Geophys. Res. Lett.*, **23**, 57–60.
- Wang, B., and Y. Wang, 1996: Temporal structure of the Southern Oscillation as revealed by waveform and wavelet analysis. *J. Climate*, **9**, 1586–1598.
- Weng, H., and K.-M. Lau, 1994: Wavelets, period doubling, and time–frequency localization with application to organization of convection over the tropical western Pacific. *J. Atmos. Sci.*, **51**, 2523–2540.
- Wright, P. B., 1984: Relationships between indices of the Southern Oscillation. *Mon. Wea. Rev.*, **112**, 1913–1919.
- , 1989: Homogenized long-period Southern Oscillation indices. *Int. J. Climatol.*, **9**, 33–54.
- Yamada, M., and K. Ohkitani, 1991: Orthonormal wavelet analysis of turbulence. *Fluid Dyn. Res.*, **8**, 101–115.