

Quantifying Persistence in ENSO

JOHN P. WEISS AND JEFFREY B. WEISS

Program in Atmospheric and Oceanic Science, University of Colorado, Boulder, Colorado

(Manuscript received 17 December 1997, in final form 7 December 1998)

ABSTRACT

The seasonal dependence of predictability in ENSO manifests itself in the so-called spring barrier found in the cyclostationary lag autocorrelations, or persistence. This work examines the statistics of persistence, with particular focus on the phase-of-year-dependent pattern found in ENSO data, the barrier. Simple time series of one sine wave produce a barrier if the frequency is a biennial cycle or one of its harmonics. Time series of two sine waves produce a barrier if one frequency is a biennial cycle or a harmonic thereof. They additionally produce a barrier if their frequencies sum to unity. Time series with continuous but narrow spectral peaks at barrier-producing frequencies produce barriers only if the phase angles vary slowly or coherently across the peaks. The shape of the barrier seen in these simple time series is used to construct a model persistence map, which is a combination of an idealized barrier and the persistence of a red-noise process. A nonlinear least squares fit of the persistence of a time series to the model persistence provides a quantitative measure of the properties of the persistence barrier in any time series. Application of the measure to the Southern Oscillation index and sea surface temperature in the NINO3 region of the equatorial Pacific indicates that the ENSO persistence barrier is a feature that is statistically distinguishable from the theoretical persistence of a red-noise process. The ENSO barrier results from phase coherency of the continuum of interannual modes near the biennial frequency. Measuring the barrier on windowed data shows that there was a weak persistence barrier from approximately 1915 to 1945, a strong barrier during the 1960s and early 1970s, and a weakening barrier in the late 1970s.

1. Introduction

The El Niño–Southern Oscillation, or ENSO, is one of the most readily recognizable forms of interannual climate variability of the coupled atmosphere–ocean system. It also has broad human impact, with effects ranging across the equatorial Pacific basin and teleconnections to much of the globe (Bjerknes 1969; Philander 1990). Consequently, there exists a large body of work examining the basic physics of ENSO (Battisti and Hirst 1989; Bjerknes 1966; Cane and Zebiak 1985; Zebiak and Cane 1987; Rasmusson and Carpenter 1982; Vallis 1988). Among this work, Troup (1965) and others (Wright 1979, 1985; Webster and Yang 1992; Webster 1995; Torrence and Webster 1998) have examined the persistence of ENSO via the Southern Oscillation index (SOI). Persistence, as defined in these works, is a form of lag autocorrelation. A map of the persistence of the SOI, similar to ones shown by Wright (1979, 1985) and Torrence and Webster (1998), is shown in Fig. 1. Note the rapid decrease in persistence during the boreal

spring. Webster and Yang (1992) termed this decrease a “predictability barrier,” though it is not yet clear how this decrease in persistence relates to the ability to forecast ENSO. There has been work suggesting that this decrease in persistence may not necessarily hinder forecasting (Chen et al. 1995). Nonetheless, this feature of ENSO persistence is intriguing. Torrence and Webster (1998) show that a similar barrier appears in the persistence of another ENSO index, the sea surface temperature in the NINO3 region of the central equatorial Pacific (Fig. 2). They suggest that such a barrier occurs as the climate system makes a transition from one state to another during the boreal spring. Xue et al. (1994) additionally show that the barrier is associated with low variance in the data during the spring.

Although there is some controversy over the meaning of the spring persistence barrier, it is nonetheless a statistical property of ENSO. As such, it warrants further investigation, as deeper understanding of persistence barriers may provide further insight into the causes and dynamics of ENSO. In addition, statistical features, even if they are not fully understood, can provide useful benchmarks for comparing models with data. Thus, in this paper we study the statistics of persistence from a primarily mathematical perspective. The results presented here demonstrate that the persistence barrier in ENSO is a statistically significant feature that arises

Corresponding author address: Dr. Jeffrey B. Weiss, Program in Atmospheric and Oceanic Sciences, Campus Box 311, University of Colorado, Boulder, CO 80309-0311.
E-mail: jweiss@colorado.edu




FIG. 1. Persistence map of the SOI from 1896–1996. Contour intervals are in steps of $\frac{1}{4}$. Solid contours denote positive values of persistence, while dotted contours denote negative values. The zero contour is dashed.

from particular features in the ENSO time series. While we do not in this paper determine the underlying dynamical mechanisms that give rise to the necessary properties of the time series, the results suggest that these mechanisms are generic and that predictability barriers should occur in general nonlinear oscillators and not just ENSO.

The appropriate class of statistics for time series with an external periodicity, such as an annual cycle, is cyclostationary statistics (Huang and North 1996). Section 2 presents a formal definition of cyclostationary statistics, and then calculates the persistence for a progression of simple but increasingly complex time series. One simple question is whether the persistence of a time series depends on phase of year. We find that some time series have a persistence that is independent of the phase of the year, others have phase dependence that resembles the persistence barrier seen in ENSO data, while yet others have a phase-dependent persistence that looks completely different than that seen in ENSO data. Our focus here will be on those phase-dependent patterns that match the ENSO data. We will refer to these patterns as barriers. In section 2 we show that a persistence barrier is due to the presence of harmonics in a time series whose frequencies satisfy certain conditions. In section 3 we develop a method to quantitatively measure the properties of a persistence barrier. In section 4a, we use this measure to calculate the properties of the barriers in the NINO3 SST and the SOI. Section 4b numerically examines the conditions for producing a barrier, the conditions described analytically in section 2. Finally, in




FIG. 2. Persistence map of the NINO3 SST from 1896–1996. The contours are plotted in the same fashion as in Fig. 1.

section 5 we use the measure of barrier properties to study the interdecadal variability of ENSO.

2. The mathematics of persistence

a. Definition

Common forms of statistical analysis assume that the data is stationary in time. Most geophysical systems, however, are characterized by periodic cycles, such as the annual cycle or the diurnal cycle. It is thus appropriate to include this external periodicity in the statistical analysis; this results in cyclostationary statistics (Huang and North 1996). Cyclostationary statistics differ from stationary statistics in that the various quantities, such as the mean, variance, etc., are periodic functions of the phase with respect to the external periodic cycle.

In this paper, the relevant external periodicity is the annual cycle, and we will measure time in units of years. For the following definitions we consider a time series $x(t)$ and denote the phase of the year by p , $0 \leq p < 1$. The cyclostationary mean is defined as

$$\langle x(t) \rangle(p) = \frac{1}{N} \sum_{y=0}^{N-1} x(y+p). \quad (1)$$

Thus, $\langle \cdot \cdot \cdot \rangle(p)$ represents an average over N years at fixed phase p . In the case of the ENSO data, Eq. (1) is merely the mean annual cycle. It will be convenient to also write the cyclostationary mean as $\langle x \rangle(t)$ where the dependence on t is understood to mean the phase of t . The cyclostationary variance is defined in terms of the departure from the annual cycle, that is, the anomaly, $x'(t) = x(t) - \langle x \rangle(t)$,

$$\sigma_y^2(p) = \frac{1}{N-1} \sum_{y=0}^{N-1} [x'(y+p)]^2. \quad (2)$$

The cyclostationary lag autocorrelation, also called persistence, is then

$$C(p, l) = \frac{\langle x'(t)x'(t-l) \rangle(p)}{\sigma_y(p)\sigma_y(p-l)}. \quad (3)$$

We will refer to p as the leading phase, or just phase if there is no confusion. We will focus our attention on lags of 1 yr or less, $l \leq 1$, and will refer to the resulting field as a persistence map. Note that p is a specific phase of the year, for example, 1435 UTC 12 February, while l is a time interval. One could, of course, define the persistence in terms of its starting and ending phases, $p_1 = p$ and $p_2 = p + l$. However, since persistence always depends on l and we are investigating when and how it depends on p , the above convention makes more sense.

b. Persistence of one sine wave

One step in understanding the meaning of a statistical analysis such as persistence is to calculate the statistics of simple processes. For example, in interpreting a power spectrum, one compares to the spectrum of a red-noise process to see if any peaks are distinguishable from random noise, and compares those peaks with those of a linear oscillator to see if they can be understood in terms of internal oscillations. In this spirit we calculate the persistence of several simple deterministic time series as well as that of red noise in order to investigate when one finds persistence barriers in simple systems. Since any complex time series can be decomposed into harmonic functions by the Fourier transform, it is reasonable to consider the persistence of simple harmonic time series. Note that since persistence is a nonlinear function of a time series, simply superposing the persistence of many pure harmonic functions does not suffice. However, proceeding through a sequence of increasingly complex time series does provide insight into the properties that a time series must possess to produce a persistence barrier. In all the cases below, persistence is calculated with respect to an annual cycle, that is, period of 1 yr, while the frequencies in the time series vary.

First, consider a single, deterministic, sinusoidal periodic oscillation,

$$x(t) = a \sin(2\pi\nu t + \phi), \quad (4)$$

where $\nu > 0$ (negative frequencies are equivalent to redefining ϕ). The calculation of the persistence map for a single sine wave is presented in appendix A; here we describe the results.

The persistence of a simple sine wave divides into three cases depending on whether 1) ν is an integer n [Eq. (A2)], 2) ν is a half-integer $(2n+1)/2$ [Eq. (A3)], or 3) neither [Eq. (A4)]. A time series with $\nu = n$ cor-




FIG. 3. Persistence map of a single sine wave, Eq. (5), with $\beta = 0.25$. The shaded area denotes a persistence of +1, while the unshaded area is a persistence of -1.

responds to an annual cycle or a harmonic of an annual cycle. Since the persistence is calculated from the departure from the annual cycle, it is undefined in this case. If ν is neither an integer nor a half-integer, then the persistence does not depend on leading phase.

The remaining case, where ν is a half-integer, corresponds to a biennial cycle or one of its harmonics. In this case, the persistence $C(p, l)$ can be written in the form

$$C(p, l) = \Theta(p - \beta)\Theta(p - \beta - l), \quad (5)$$

where $\beta \equiv -\phi/2\pi\nu$, $\Theta(x) = (-1)^{\text{Int}(x)}$ is a periodic step function (i.e., a square wave) and $\text{Int}(x)$ denotes the integer part of x . Thus, C depends on both leading phase and lag and consists of two discrete steps: one step independent of l at phase $p = \beta$, and one at $l = p - \beta$. A contour plot of Eq. (5) for $\nu = 1/2$ appears in Fig. 3. Note the similarities to the persistence maps of the SOI and NINO3 sea surface temperature, Figs. 1 and 2. The parameter β determines the position of the transition in the map from positive to negative persistence. Note that β is also the time at which $x(t)$ crosses zero.

We have also examined the persistence of a single harmonic of the biennial cycle, $\nu = n/2$, $n \neq 1$. The persistence still depends on leading phase and contains discrete steps; however, the steps now repeat multiple times in both p and l .

c. Persistence of two sine waves

We now turn to a time series composed of two sine waves:

$$x(t) = a_1 \sin(2\pi\nu_1 t + \phi_1) + a_2 \sin(2\pi\nu_2 t + \phi_2), \quad (6)$$

where ν_1 and ν_2 are positive numbers. Appendix B contains the derivations and results of the persistence of (6). Again, the persistence breaks into three cases: 1) ν_1 or ν_2 are integers, 2) neither ν_1 nor ν_2 are integers and at least one is a half-integer, and 3) ν_1 and ν_2 are neither integers nor half-integer.

When either ν_1 or ν_2 is an integer, that Fourier mode corresponds to an annual cycle. Computing the anomaly removes that harmonic of the time series and the persistence is identical to that of the single remaining sine wave.

For the case where at least one of the frequencies is an odd half-integer, one might expect that the persistence resembles that of a single biennial sine wave [i.e., Eq. (5)]. Such a resemblance does indeed exist, as seen in Fig. 4. Only when both frequencies are half-integers, however, does $C(p, l)$ have discrete steps. When only one of the two modes is a biennial cycle, the drop in correlation is continuous, as for the ENSO data.

We have additionally explored the behavior of the persistence $C(p, l)$ for values of $a_1, a_2, \phi_1,$ and ϕ_2 different than those used in Fig. 4. The phases ϕ_1 and ϕ_2 determine the position of any leading-phase-dependent features. When $C(p, l)$ is continuous, (i.e., only one harmonic is biennial in nature), the values of a_1 and a_2 control the strength of the gradients of the continuous features. For example, decreasing $|a_1 - a_2|$ would narrow the transition region between positive and negative persistence. Furthermore, if $a_1 \gg a_2$, then $x(t) \approx a_1 \cos(2\pi\nu_1 t + \phi_1)$ and the persistence changes accordingly.

Figure 4 illustrates forms of phase dependence other than that found in the persistence of the SOI and NINO3 SST. In this work, we shall use the term “barrier” to denote those phase-dependent patterns in the persistence map that are qualitatively similar to the pattern seen in the SOI and NINO3 SST.

In the final case, where ν_1 and ν_2 are neither integers nor half-integers, one might expect the persistence to be independent of p . Such behavior would be consistent with the corresponding case for the single sine wave. Surprisingly, $C(p, l)$ is only a function of lag if $\nu_1 \pm \nu_2 \neq m$, where m is an integer. Figure 5 contains plots of the persistence for the $\nu_1 \pm \nu_2 = m$ case. (The mathematics behind these figures can be found in appendix B.) Note how different values of m correspond to different forms of leading-phase dependence.

The discussion of Eq. (B5) in appendix B reveals some additional characteristics of the persistence map for the case $\nu_1 \pm \nu_2 = m$. The values of $\nu_1, \nu_2, m,$ and whether the two frequencies sum to or differ by m all control the shape of any leading-phase-dependent features. The phases ϕ_1 and ϕ_2 determine the position of the features. Finally, the difference of the Fourier amplitudes, a_1 and a_2 , controls the gradients of the transition regions between different features in the persis-

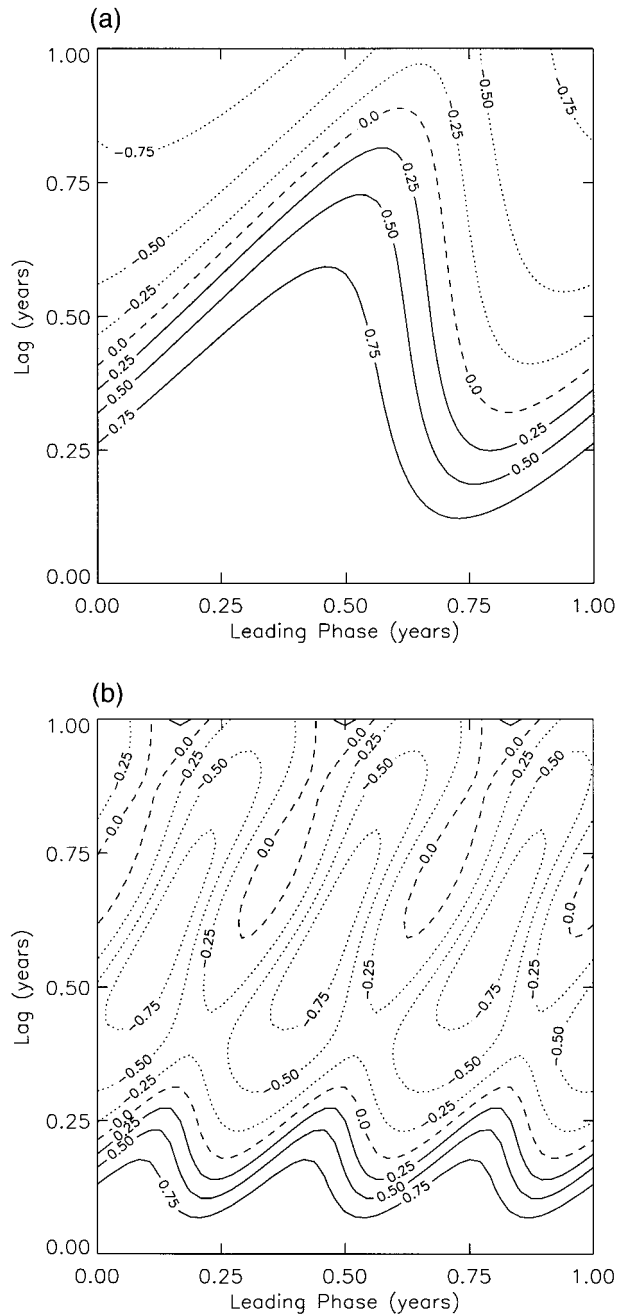


FIG. 4. Persistence maps of time series composed of two harmonics, Eq. (6), with $a_1 = 1.0, a_2 = 0.75,$ and $\phi_1 = \pi/3, \phi_2 = \pi/2,$ and with a biennial cycle present. Contour intervals are in steps of $1/4$, with the contour lines drawn as described in Fig. 1. (a) $\nu_1 = 1/2$ and $\nu_2 = 3/10$; (b) $\nu_1 = 1/5$, and $\nu_2 = 3/2$.

tence map. Smaller $|a_1 - a_2|$ correspond to stronger gradients, with $a_1 = a_2$ yielding a step function.

The special case of $\nu_1 \pm \nu_2 =$ an integer has implications for the persistence of time series beyond the sum of two sine waves. We will therefore denote any pair of harmonics with frequencies ν_i and ν_j that satisfy $\nu_i \pm \nu_j = m$ as m^\pm complementary. From Fig. 5a, one

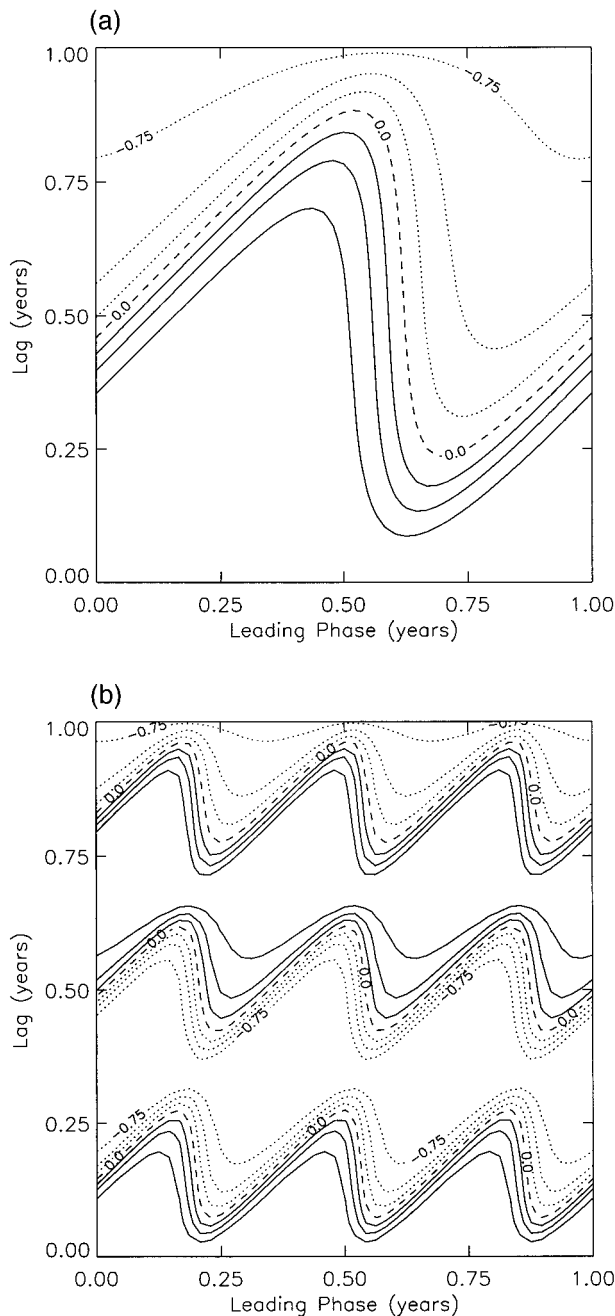


FIG. 5. Persistence maps of time series composed of two harmonics, Eq. (6), with $a_1 = 1.0$, $a_2 = 0.5$, $\phi_1 = \pi/3$, $\phi_2 = \pi/2$, and $\nu_1 \pm \nu_2 = m \in \mathbb{Z}$. Contours are drawn as described in Fig. 4. (a) $\nu_1 = 2/5$, $\nu_2 = 3/5$, $\nu_1 + \nu_2 = 1$; (b) $\nu_1 = 7/5$, $\nu_2 = 8/5$, $\nu_1 + \nu_2 = 3$.

sees that $\nu_i + \nu_j = 1$ results in a barrier-like feature, while other m^\pm -complementary harmonics produce other forms of phase dependence. We will refer to harmonics with this relationship, $\nu_i + \nu_j = 1$, as unit complementary. Note that both frequencies, ν_i and ν_j , in a unit-complementary pair must be less than 1, since both frequencies are also positive.

Summarizing the results of the previous two sections,

Table 1 describes those cases where the persistence contains some form of phase dependence. Persistence takes the form of a barrier for a single biennial mode, the sum of a biennial and nonbiennial mode, or for the sum of a unit-complementary pair of harmonics. Other forms of phase dependence occur for a biennial harmonic or other m^\pm -complementary pairs.

d. Persistence of a general time series

A general time series is described by a continuous Fourier transform,

$$x(t) = \int_0^\infty da(\nu) \sin[2\pi\nu t + \phi(\nu)], \quad (7)$$

and the amplitudes and phase angles are now functions of the frequency. For the general case, calculating the persistence in closed analytic form appears intractable.

There are a few special cases, however, where the persistence can be calculated. Consider a time series, $x(t)$, whose spectrum consists of either a single biennial peak or two unit-complementary peaks. Assume the peaks are sufficiently narrow and, in the case of two unit-complementary peaks, do not overlap. The persistence of x can then be calculated in certain limits. The derivation is lengthy and not straightforward. One can find its more pertinent points in appendix C. Note that the calculation requires the power spectrum of $x(t)$ to have a certain distribution, not x itself.

In these cases the persistence has two different limits. If the phase angle $\phi(\nu)$ is rapidly varying over the width of a peak, then the persistence depends only on lag [Eq. (C3)]. This holds not only for the single biennial peak, but for a pair of unit-complementary peaks as well. It can be generalized to multiple sets of unit-complementary harmonics: if the phase varies rapidly over all of the peaks, then the resulting persistence map is independent of the phase of year.

On the other hand, in the limit of slowly varying $\phi(\nu)$, the persistence does contain a barrier [Eqs. (C4) and (C5)]. Figure 6 contains a plot of the persistence in this limit. It shows a strong similarity to the single mode cases of sections 2b and 2c. For both a single biennial peak and two unit-complementary peaks, broadening the peaks broadens the transition region.

Note that the assumptions used in appendix C to compute the persistence for these cases do not strictly require slowly varying $\phi(\nu)$. The assumptions are also satisfied by $O(1)$ slopes in the Fourier phase, provided the slopes are similar at the two peaks. Thus, Eqs. (C4) and (C5) are valid when the phases of the two spectral peaks vary coherently.

We can generalize these results by the following hypothesis. A time series with a continuous power spectrum will contain a barrier provided that 1) barrier-generating modes, that is, a biennial mode or pairs of unit-complementary modes, have significant amounts of

TABLE 1. Summary of phase-dependent persistence maps for time series of a single sine wave with frequency ν [Eq. (4)] and the sum of two sine waves with frequencies ν_1 and ν_2 [Eq. (6)].

Persistence pattern	ν	ν_1, ν_2
Barrier	Biennial [Eq. (5)]	Biennial + nonbiennial [Eq. (B3)] Unit complementary $\nu_1 + \nu_2 = 1$ [Eq. (B5)]
Other phase dependence	Biennial harmonic (Eq. (5))	Biennial + any biennial harmonic [Eq. (A5)] Two biennial harmonics [Eq. (A5)] m^\pm complementary with $\nu_1 + \nu_2 \neq 1$ [Eq. (B5)]

power; and 2) the Fourier phases of the barrier-generating modes vary coherently. The first requirement implies that the time series has strong interannual variability, while the second requirement is a particular form of phase locking. We shall examine whether this hypothesis holds for the ENSO datasets in section 4b.

e. Persistence of a stochastic time series

When one is asking whether a statistical feature is significant, one must compare it with a null hypothesis. Typically, in the geophysical sciences the null hypothesis is a random noise process with a red power spectrum. Since a red-noise process is stationary, not just cyclostationary, its persistence is independent of leading phase and is equal to its lagged autocorrelation.

The red-noise process we consider is the first-order Markov process known as the Ornstein–Uhlenbeck process,

$$\dot{x} = -\frac{x}{\tau} + \xi(t), \quad (8)$$

where $\xi(t)$ is Gaussian white noise with zero mean and




FIG. 6. Persistence map calculated for a time series whose spectrum contains two narrow peaks with coherently varying phase angles, Eq. (C4), with $\alpha = 2$, $\hat{\nu} = -0.2$, $\beta = 0.5$, $\Delta = 0.1$, and $\epsilon_2 = \epsilon_4 = 0.05$.

unit variance. The noise process is described by the parameter τ , the timescale of the deterministic part. Using standard techniques of stochastic calculus (Gardiner 1990), it is straightforward to show that the ensemble average persistence is simply

$$\langle C(p, l) \rangle_{\text{ensemble}} = e^{-l/\tau}. \quad (9)$$

The parameter τ , therefore, is also the decorrelation time for the persistence.

3. A measure for persistence barriers

a. Definition and properties

In order to further study the barrier in a time series one needs to quantify its properties. Once a statistical feature is quantified, one can ask about its statistical significance, its temporal variability, and how it is affected by changing the dynamics of the system, such as adding noise, increasing resolution, or coupling to another system. Further, one can quantify how well a numerical model reproduces the feature seen in the data. Thus, in this section, we develop a means to quantitatively measure the persistence barrier in a time series.

A persistence map is a two-dimensional field. Creating a measure of a persistence barrier necessarily involves some form of pattern recognition, one that reduces said field down to one or a few numbers. Typically, pattern recognition requires comparing data with some idealization of the pattern one wants to find. Thus, we construct a theoretical persistence barrier, $T(p, l)$, which can be used to measure the barrier in the persistence map, $C(p, l)$, from a time series $x(t)$. The model barrier should be described by a small number of parameters. To determine the properties of the barrier in x , one then performs a least squares fit of T to C . The values of the parameters at the best fit provide the barrier measure.

From the results of the previous section, it is clear that an idealized persistence barrier should consist of two jumps in persistence, one independent of lag at a phase $p = p_b$, and one at lag $l = p - p_b$. We shall call p_b the barrier position. Further, the jumps in persistence can be either sharp or broad. Thus, the barrier should be characterized by a barrier width, Δp . We also wish to allow the possibility that the persistence map does not have a barrier at all but rather arises from a red-noise process. Thus, we introduce an amplitude, A , that measures the strength of the barrier. Finally, since we

include a red-noise component, we need to include the decorrelation time of the stochastic process, τ . Thus, the simplest barrier that suits our needs requires four parameters: an amplitude, A ; a width, Δp ; a position, p_b ; and a decorrelation time, τ . We note that the model barrier we construct is chosen merely because it is simply described and has the appropriate shape.

The model persistence map is thus the sum of two terms:

$$T(p, l) = AB(p, l) + (1 - A)e^{-l/\tau}, \quad (10)$$

a barrier, $B(p, l)$, with amplitude A , and the persistence of the Ornstein–Uhlenbeck stochastic process [Eq. (8)] with amplitude $1 - A$. To complete the model we must define $B(p, l)$. One could define B using the persistence of the two-mode case with m^\pm -complementary frequencies [Eq. (B5)]. However, the parameters in Eq. (B5) are difficult to interpret, and fitting to the equation is computationally prohibitive. Instead, we use a model barrier that has a simpler form.

The jumps in the barrier should be smooth transitions from -1 to $+1$ at a particular location. A simple function that achieves this is $\tanh(p/\Delta p)$, where Δp measures the width of the transition region. Since persistence must be periodic in phase, we construct the function

$$h(p, \Delta p) = \sum_{n=-\infty}^{\infty} \tanh\left(\frac{p+n}{\Delta p}\right), \quad (11)$$

which has transitions at $p, p \pm 1$ yr, $p \pm 2$ yr, etc. Since we want two transitions, one at $l = p - p_b$ and one with opposite sign at $p = p_b$, we choose

$$B(p, l) = h(p - p_b - l, \Delta p) - h(p - p_b, \Delta p) + 1, \quad (12)$$

where the final term places B in $[-1, +1]$. Computationally, we truncate the sum in Eq. (11) at $20/\Delta p$, resulting in errors on the order of machine precision. We now have a theoretical barrier, T , that is completely described by a vector of four parameters, $\mathbf{a} = (A, \tau, \Delta p, p_b)$.

Because T models a correlation function, it must satisfy the conditions $-1 \leq T(p, l) \leq 1$ and $T(p, l) = 1$ at $l = 0$. This leads to the following requirements on the parameters: $0 \leq A \leq 1$, $\Delta p \geq 0$, and $\tau \geq 0$. By construction, T is periodic in p_b with unit period. The parameters p_b , Δp , and τ all have units of time, while A is unitless. The mean of B is

$$\overline{B} = \int_0^1 dl \int_0^1 dp B(p, l) = 0, \quad (13)$$

and thus the mean of T , \overline{T} , is given by the mean of the stochastic part, $\overline{T} = (1 - A)\tau(1 - e^{-1/\tau})$.

To interpret the parameters, it is useful to first investigate the asymptotic behavior of T for large τ and Δp . We only consider lags restricted to $0 \leq l \leq 1$. For large τ , $T \sim AB + (1 - A) + O(\tau^{-1})$; large decorrelation

times thus result in the stochastic portion providing a constant background to the persistence map. For large Δp , $B \sim (1 - 2l) + O(\Delta p^{-3})$, and the barrier reduces to a phase-independent, linear decay with lag. Numerically, these asymptotic results prove to be good approximations when $\tau > 10$ yr or $\Delta p > 1$ yr.

The model barrier T contains a rich variety of behaviors, reflected in the physical interpretation of its parameters. In particular, the three parameters A , τ , and Δp occasionally have different implications depending on their joint values. We will now discuss some of the different behavior regimes of the model and the values of A , τ , and Δp associated with them.

When A is nearly one, that is, $1 - A \ll 1$, the persistence describes a deterministic process. The model contains regions of high positive and negative correlation, separated by a transition region whose width is determined by Δp . Thus, any loss of correlation is temporary and the barrier represents only a transition from correlation to anticorrelation instead of a loss of predictability. An example is shown in Fig. 7a. In the limit $A \rightarrow 1$, the decorrelation time τ has no meaning.

If the barrier amplitude is small, $A \ll 1$, the persistence map is that of a random process, that is, a decaying exponential, modified by a small phase dependence. In this regime, the decorrelation time τ determines how rapidly correlation is lost. In the limit $A \rightarrow 0$, the barrier position p_b and barrier width Δp have no meaning.

For large barrier width, $\Delta p \geq 1$, the persistence becomes independent of phase. The values of A and τ together determine the decorrelation rate. In this regime, the deterministic part of the model has a linear decorrelation, which combines with the exponential decay from the stochastic portion.

When the decorrelation time becomes large, $\tau \gg 1$, the stochastic component of the model is approximately constant. The full model thus has regions containing high positive correlation and regions containing correlation with a value of $1 - 2A$. Thus, for appropriate values of A , the model varies from regions of high correlation to regions of correlation near zero. The transition region separating the two has a width determined by Δp . One can also think of Δp as the rate of decorrelation with respect to phase. Thus larger values of Δp indicate weaker barriers. An example of this case is shown in Fig. 7b for $A = 0.6$.

b. Computation

The parameters describing the persistence map, $C(p, l)$, of a time series can now be found by fitting T to C . For a time series with finite resolution, C will be only available on a discrete grid of points in (p, l) . The fit is accomplished by minimizing the error function,

$$E^2(\mathbf{a}) = \sum_{p,l} [T(p, l; \mathbf{a}) - C(p, l)]^2. \quad (14)$$

This nonlinear fit suffers from the usual problem of




FIG. 7. The quantity $T(p, l)$ at various parameter values. For all plots, $p_b = \frac{1}{4}$ yr. (a) $A = 1$, $\Delta p = 0.06$ yr, $\tau = 0.25$ yr; (b) $A = 0.6$, $\Delta p = \frac{1}{4}$ yr, $\tau = 10$ yr.

multiple local minima in E^2 . To handle this difficulty, we use a custom genetic algorithm (Goldberg 1989) to locate the neighborhood near the global minimum, and then use a nonlinear least squares fit to find the actual minimum. The genetic algorithm utilizes some standard techniques such as elitism; however, it also encodes the parameters of T as a vector of real numbers while utilizing a high mutation rate (Haupt and Haupt 1998).

The least squares fit uses a Levenberg–Marquardt algorithm implemented by the routine “lmdr.f” from the Minpack library (Moré et al. 1980).

The uncertainty in the fitted parameters is found using standard techniques (Ryan 1997; Press et al. 1992). Denote the value of the parameters that minimizes E^2 as \mathbf{a}_0 . Then the standard estimate of the confidence interval on the k th parameter at the q th percentile is

$$\delta a_k(q) = \left[\frac{E^2(\mathbf{a}_0) D_{kk}^{-1} \chi_1^2(q)}{N} \right]^{1/2}, \quad (15)$$

where $D_{ij} \equiv (\partial_{a_i} \partial_{a_j} E^2(\mathbf{a}_0))/2$ is the curvature matrix of the error function at \mathbf{a}_0 , D_{kk}^{-1} is the k th diagonal element of the inverse of the curvature matrix, N is the number of data points in the map, and $\chi_1^2(q)$ is the percentile value of a one degree of freedom chi-square distributed variable at the q th percentile. For $q = 68.3\%$, 95.4% , or 99.7% , $\chi_1^2(q)$ is simply 1, 4, or 9, respectively.

There is one complication to calculating the confidence intervals. The above equation assumes that the local curvature, obtained from the second derivative of E^2 , is a good approximation to the average curvature over the width δa_k . In certain cases, however, the E^2 surface is much flatter at the minimum than farther away. The curvature matrix obtained from the second derivative is then ill conditioned or singular. An example is the case where the width of the barrier is much less than the resolution of the data. In this case, infinitesimal changes in the barrier position barely affect the error and the associated derivative is zero. Changes in position on the order of the data resolution, however, do cause significant changes in the error. Thus, rather than use analytic derivatives of E^2 , which are necessarily local, we use finite difference derivatives to compute an average curvature. The resulting curvature matrix only becomes ill conditioned when A is near 0 or 1, where some of the parameters become ill defined. For this special case, we directly assign a large value to the δa_k of the ill-defined parameters. We then calculate the confidence intervals of the remaining parameters by inverting the submatrix of the curvature matrix that corresponds to the well-defined parameters. We can then calculate the confidence as usual.

While one can always find a fit that minimizes the error, we would like to know whether the fit is any good. One measure of how well two datasets are linearly correlated is the familiar correlation coefficient. Here, however, T is a nonlinear fit and the usual correlation coefficient is a rather poor goodness-of-fit criterion. Kvålseth (1985) discusses the issues surrounding several forms of nonlinear correlation coefficient, R^2 , and recommends using a particular form in most cases. Applied to persistence maps, this form is

$$R^2 = 1 - \frac{E^2(\mathbf{a}_0)}{\sigma_{C(p,l)}^2}, \quad (16)$$

where \mathbf{a}_0 represents the best-fit parameters of the model,

TABLE 2. Results of the various Monte Carlos performed on the NINO3 SST. Entries are either the median values plus or minus the 68% confidence interval of the distribution for the Monte Carlo, or the actual measure plus or minus the estimated 68% confidence interval. The labels for each type of Monte Carlo are explained in the text.

Type	A	p_b	Δp	τ
Data	0.312 ± 0.018	0.258 ± 0.005	0.137 ± 0.011	1.072 ± 0.081
OU	0.048 ± 0.037	0.501 ± 0.337	0.136 ± 0.131	1.113 ± 0.219
FPS	0.263 ± 0.012	0.509 ± 0.337	0.345 ± 0.086	0.793 ± 0.028
UCPS	0.268 ± 0.013	0.492 ± 0.354	0.349 ± 0.087	0.803 ± 0.029
NCPS	0.312 ± 0.021	0.279 ± 0.011	0.148 ± 0.017	1.029 ± 0.078

and $\sigma_{C(p,l)}^2$ is the variance of the persistence map being measured. Note that this is the standard variance, computed over all points in $C(p, l)$.

In general, there may be cases where Eq. (16) returns a value of $R^2 < 0$. This situation is due either to a complete lack of fit between the data and the model, or to outliers in the data. Data outliers are clearly not an issue for the fit of T to C , two bounded quantities. Other studies, however, may not be so fortunate. Those researchers requiring a nonlinear correlation coefficient insensitive to outliers should consult Kvålseth for a more appropriate definition of R^2 .

4. ENSO barriers

a. The data

We now use the aforementioned technique to measure the properties of the persistence barrier in ENSO. The two sets of data we use are the SOI, and the SST from the NINO3 region of the equatorial Pacific, 5°N – 5°S by 150° – 90°W . Both datasets contain monthly mean values for the period 1896–1996. The NINO3 SST data from 1950 to the present is available from the National Oceanic and Atmospheric Administration Climate Prediction Center. The historical NINO3 SST data, from 1896–1949, is based on a spatial average of data from the Comprehensive Ocean–Atmosphere Data Set SST analysis (Oort et al. 1987). The SOI is a standardized difference in the sea level pressure between Tahiti, at 17.6°S , 149.6°W , and Darwin, Australia, at 12.4°S , 130.9°E . The pressure at each location has been standardized by removing its cyclostationary mean and normalizing by its overall standard deviation. The SOI is then computed by taking the difference in the two standardized pressures and normalizing that difference by its overall standard deviation. To fill the gaps that exist in the Tahiti sea level pressure record before 1950, Torrence and Webster (1998) perform a least squares re-

gression using sea level pressure and sea surface temperature from other locations in the equatorial Pacific. The contiguous SOI data for 1896–1950 produced by this analysis is used in this work.

The results of fitting the model persistence barrier to the NINO3 SST and the SOI are shown in Tables 2 and 3, respectively, in the row labeled “data.” The reported uncertainty is the 68.3% confidence interval estimated using Eq. (15). The NINO3 SST barrier has a strength, A , roughly three times larger than that of the SOI barrier, is about four times wider, and has a longer stochastic decorrelation time. The barriers in both datasets have positions at approximately the same time of year, middle to late March. The 68.3% confidence intervals are at most about $1/10$ of the parameter values, indicating that the parameters are relatively well defined. The sole exception is the SOI barrier width whose uncertainty is relatively large because the barrier is so narrow. Since A is bounded away from zero, both the NINO3 SST and the SOI have statistically significant barriers. The shorter decorrelation time in the SOI is understood in terms of the faster typical timescale for atmospheric fluctuations, and the fact that the atmosphere is often thought of as the more random component of the coupled system.

For the NINO3 SST and SOI barriers, the goodness-of-fit descriptors (defined in section 3b) are $R^2 = 0.966$ and 0.841 , respectively. Evidently, it is not unusual for R^2 to have values greater than 0.9 (Kvålseth 1985). Thus, the persistence barrier of the SOI does not match the model too well, indicating the possible presence of other forms of phase dependence in its persistence map.

Let us now investigate whether the measured persistence barriers can be explained by random fluctuations due to the finite size of the data sampling. This is done by integrating the Ornstein–Uhlenbeck (OU) process of Eq. (8). An infinitely long OU time series produces a persistence map with pure exponential decay. The

TABLE 3. Results of the various Monte Carlos performed on the SOI. The notation is identical to that used in Table 2.

Type	A	p_b	Δp	τ
Data	0.095 ± 0.012	0.275 ± 0.012	0.031 ± 0.029	0.398 ± 0.016
OU	0.049 ± 0.036	0.500 ± 0.336	0.072 ± 0.091	0.418 ± 0.060
FPS	0.037 ± 0.018	0.503 ± 0.338	0.007 ± 0.042	0.357 ± 0.010
UCPS	0.038 ± 0.019	0.460 ± 0.396	0.007 ± 0.042	0.361 ± 0.010
NCPS	0.084 ± 0.013	0.313 ± 0.032	0.068 ± 0.046	0.381 ± 0.007




FIG. 8. Power spectrum of the NINO3 SST with the mean annual cycle removed. The power of each mode is normalized by the total power of the time series. The dotted line indicates the biennial harmonic.

ENSO data, however, is not infinite, but is approximately 100 years of monthly mean data. To investigate the uncertainty associated with finite-sized time series, we integrated 10 000 realizations of the OU process for 100 yr, using appropriate techniques for the numerical integration of stochastic differential equations (Gardiner 1990). For each realization, we then calculated the best-fit parameters for the persistence map obtained from monthly mean data. Two such experiments were performed: one with $\tau = 1$ yr to compare with the NINO3 SST, and one with $\tau = 0.4$ yr to compare with the SOI. The rows labeled “OU” in Tables 2 and 3 show the median best-fit parameters and the confidence intervals that account for 68.3% of the data from the 10 000 realizations.

This experiment shows that a red-noise time series with length similar to the ENSO datasets has a typical barrier strength of 0.05, independent of the other parameters. In both runs, $p_b \approx \frac{1}{2}$ with a confidence interval of about $\frac{1}{3}$, similar to the statistics of a flat, uniform distribution. The width Δp has a confidence interval that encloses zero. Regardless, neither p_b nor Δp are well defined at the small value of A found in these experiments. Lastly, to 68% confidence, τ is at the value used in the given Monte Carlo experiment.

Let us now compare the Monte Carlo experiments performed on the OU process with the ENSO data. For the NINO3 SST barrier, such a comparison shows that the barrier strength is statistically different from that of the OU experiment. The NINO3 SST barrier thus cannot be explained by the finite length of the data record. In




FIG. 9. Power spectrum of the SOI, plotted as in Fig. 8.

contrast, the confidence intervals of the OU experiment and the SOI barrier strength do overlap, indicating that there is some possibility that the SOI barrier is explained by a short red-noise time series.

The goodness of fit for both Monte Carlo experiments is $R^2 \approx 0.99$ with a 68% confidence on R^2 of $\approx O(10^{-3})$. Therefore, the results of the OU experiment match the stochastic portion of the barrier model very well. Recall, however, that $R^2 = 0.841$ for the SOI. Thus, the overlap between the measures of the SOI and the OU experiment is most likely due to the poorer fit between the SOI persistence and the barrier model.

b. Barrier-generating harmonics

In section 2, we examined the persistence of several simple time series and hypothesized that a time series must satisfy certain criteria for it to produce a persistence barrier. First, a time series must either contain a dominant biennial mode, that is, $\nu = \frac{1}{2} \text{ yr}^{-1}$, or it must have dominant interannual modes that form unit-complementary pairs, that is, $\nu_1 + \nu_2 = 1 \text{ yr}^{-1}$. Second, the phase angles of the unit-complementary and biennial modes must be slowly varying or mutually coherent. In this section, we will explore whether this hypothesis is true for the ENSO data.

Figures 8 and 9 contain the power spectra of the anomalous NINO3 SST and of the SOI, respectively. The dotted line marks the biennial mode, $\nu = \frac{1}{2} \text{ yr}^{-1}$. There are several features to note. Both spectra display a broad peak in the range of 0.1–0.5 yr^{-1} (periods of 2–10 yr), which is the usual timescale for ENSO. The spectra do not have narrow peaks at the biennial mode, indicating that the barrier is not due to a single biennial

peak. Finally, intra-annual modes, that is, those with frequencies greater than 1 yr^{-1} , are typically an order of magnitude weaker in power than the dominant interannual modes.

To test which modes most strongly affect the barrier, we perform three filtering experiments on the ENSO data. Each experiment has a similar form: after calculating the Fourier transform of the data, the amplitudes of certain harmonics are set to zero, the modified transform is then inverted, the persistence is calculated, and the barrier is measured. The first experiment asks whether the most powerful unit-complementary modes produce the barrier: first the mode with the largest power is selected, its unit complement is located, both modes are removed, and the process is repeated. The second experiment asks whether unit-complementary modes nearest the biennial mode are most important: all harmonics between the frequencies $\frac{1}{2} - d\nu \text{ yr}^{-1}$ and $\frac{1}{2} + d\nu \text{ yr}^{-1}$ are filtered and the barrier is measured as a function of $d\nu$. The third experiment investigates the effects of the non-barrier-producing modes: all intra-annual harmonics, $\nu \geq 1 \text{ yr}^{-1}$, are removed, leaving only the interannual harmonics, all of which form unit-complementary pairs.

The five panels of Fig. 10 summarize the results of the three filtering experiments on the NINO3 SST. The solid lines indicate the experiment where unit-complementary pairs are removed in order of decreasing power. The dashed lines display the results of filtering out modes as a function of distance from $\nu = \frac{1}{2} \text{ yr}^{-1}$. The dotted line in Fig. 10a denotes the percentage of the total spectral power contained in the harmonics with $\nu \geq 1 \text{ yr}^{-1}$, that is, the non-unit-complementary harmonics. Thus, the unit-complementary modes of the NINO3 SST contain 85% of the power.

Figures 10b–e show the change in the barrier measures of the NINO3 SST as a function of power lost in the filtering. The bars labeled “error” indicate the 68% confidence interval averaged over all plotted data points. Values of A , p_b , and Δp with confidence intervals below 0 and larger than 1 are not plotted. The barrier strength A decreases as the power in the unit-complementary modes decreases. The other measures also change as power is lost, with the decorrelation time τ decreasing, and the barrier width and position, Δp and p_b , changing more and more drastically as A nears zero (we have truncated the plots of Δp and p_b at $A = 0$, where the latter two measures are ill defined).

Most striking is the difference in behavior of the two different methods of filtering out harmonics, especially on A and τ . The barrier strength decreases much more rapidly when the harmonics are removed as a function of distance from the biennial mode, while the converse is true for τ . Compare these results to those for the case in which all intra-annual harmonics are filtered out, indicated in the plots by the diamond: the barrier strength increases to $\frac{1}{2}$, accompanied by an increase to large τ (because $\tau \gg 1 \text{ yr}$, the diamond is absent from Fig.

10c), a slight increase in Δp , and very little change in the barrier position.

The drastic increase in τ , seen when the intra-annual modes are removed from the NINO3 SST, reveals a quirk of the fitting method described in section 3b. When the strength becomes large, the stochastic portion of the barrier model is small, and its parameter τ is poorly characterized. Further, if τ is also large, then the stochastic term becomes a small constant independent of τ . In practice, we find that when $A \geq 0.5$, the fit algorithm often seeks large decorrelation times, with τ greater than 10^6 yr .

The selective filtering experiments performed on the NINO3 SST indicate two results. First, the intra-annual modes act to weaken the persistence barrier. This is perhaps not surprising: since the harmonics with $\nu \geq 1 \text{ yr}^{-1}$ cannot form unit-complementary pairs they cannot produce a barrier themselves, and since they are fast modes they act somewhat as random noise and weaken the barrier. Second, the unit-complementary modes that have the strongest effect on the barrier strength are those nearest the biennial mode, not those with the most power.

The same three selective filtering experiments were performed on the SOI. In the SOI, the unit-complementary modes contain only 60% of the power. The behavior of the measures of the SOI barrier under the filtering experiments is similar to that of the NINO3 SST. Removing modes via distance from $\nu = \frac{1}{2} \text{ yr}^{-1}$ has a strong effect on A (and therefore on Δp and p_b , as well), stronger than that seen for the NINO3 SST. The decorrelation time remains fairly constant until most of the power in the unit-complementary modes is gone. Filtering in order of decreasing power has a less pronounced effect; A varies about its initial value until about 50% of the power is removed, when it drops to zero. This again indicates that the barrier is not produced by the most powerful modes, but rather by barrier-producing modes near the biennial mode. The decorrelation time decreases monotonically, while the width and position, as appears typical, follow the behavior of the barrier strength. Filtering the intra-annual harmonics strengthens the barrier to $A = \frac{1}{2}$, increases τ to large values, widens the barrier greatly, with $\Delta p = 0.3 \text{ yr}$, but has little effect on the barrier position.

There are several similarities between the results of the selective filtering experiments for the SOI and those for the NINO3 SST. Removing the non-unit-complementary harmonics acts to strengthen and widen the barrier. It also increases the decorrelation time τ to a value larger than 10 yr , where the stochastic portion of the model is approximately constant. A second similarity is the more rapid decrease in barrier strength when the near-biennial modes, rather than the most powerful modes, are removed. However, the barrier strength in the SOI falls off more rapidly than in the NINO3 SST, indicating less phase coherence in the SOI.




FIG. 10. Plot of barrier measures for the NINO3 SST vs fractional power remaining after filtering particular unit-complementary modes. The solid lines describe filtering in order of decreasing power. The dashed line indicates filtering via distance from $\nu = \frac{1}{2} \text{ yr}^{-1}$. The dotted line and diamonds display filtering out all non-unit-complementary harmonics. The bars labeled error indicate a typical value of the 68% confidence interval. (a) The fraction of power remaining vs the filtered harmonics, (b) the strength A , (c) the decorrelation time τ , (d) the barrier width Δp , and (e) the barrier position p_b .

c. Dependence on Fourier phase angles

We next examine the effect of the Fourier phase angles of the unit-complementary modes of the barrier. While the amplitude of a Fourier mode describes the power in a given mode, the corresponding phase angle describes when a mode changes sign. If the dominant modes have coherent phases, their sum will approach zero in a periodic fashion at a point described by the phase information. The work of Xue et al. (1994) showed that the persistence barrier in ENSO data is related to low cyclostationary variance. Their work thus supports a connection between coherent phases and the persistence barrier.

We begin our exploration of the relationship between phase angles and persistence barriers with a phase-scrambling Monte Carlo experiment. To phase scramble a time series, one first performs a Fourier transform on the data. One then replaces the phase angles with a realization of uniformly distributed independent random numbers. Inverting this modified Fourier transform completes the phase scrambling. Using 1000 realizations of phase-scrambled data, we calculate the persistence map and measures; the results are in Tables 2 and 3 in the row labeled "FPS." Because the distributions of the measures are not symmetric, we report the median of the distribution and the 68% confidence interval about the median.

Phase scrambling the NINO3 SST alters all four of the measures (Table 2). The position p_b is comparable to that found in the Ornstein-Uhlenbeck Monte Carlo experiment, indicating a uniform distribution of p_b . The goodness of fit R^2 (not reported in Tables 2 and 3), differs little from the values reported in section 4a for the unaltered datasets.

For the other measures, the confidence intervals overlap neither those of the OU experiment nor the original data. There is a decrease in the barrier strength A and a doubling of the width Δp compared to the unaltered data. Additionally, the decorrelation time τ decreases. Both the increase in width and drop in strength suggest a weakening of the barrier.

Phase scrambling the SOI renders the barrier statistically indistinguishable from a stochastic process. As seen in Table 3, the median values of A and τ are smaller than those of the OU experiments, while Δp is statistically close to zero for both experiments. The confidence intervals of the strength and decorrelation time are narrow for FPS, but do overlap those of the OU experiment. Additionally, the confidence intervals of the strength and decorrelation time do not overlap those of the unaltered data. The measures of the original data are therefore statistically different from the results of phase scrambling.

The phase-scrambling Monte Carlo experiments on the NINO3 SST and SOI reveal the importance of phase coherency to the generation of persistence barriers. In a time series with a moderately strong barrier, for ex-

ample, the NINO3 SST, removing the phase coherency weakens and broadens the barrier. When phase coherency is removed from a time series with a weak barrier, such as the SOI, the persistence becomes statistically identical to that of red noise.

As an additional test, we perform two modified phase-scrambling Monte Carlo experiments. Instead of replacing all phase angles with random numbers, we replace only a specific set of phase angles. The rows labeled "UCPS" in Tables 2 and 3 are the results of randomizing only the phases of the unit-complementary modes. The complementary experiment, where one phase scrambles all modes that are not unit complementary, is labeled "NCPS" in the tables.

The results of these two Monte Carlo experiments confirm the role of unit-complementary harmonics in generating a persistence barrier. The measures found in the UCPS experiment for both datasets are nearly identical to those of the FPS. Similarly, the measures of the NCPS experiment and those of the original data also overlap. The sole exception is the barrier position. In both datasets, p_b for the NCPS is later in the year than for the original barrier position.

5. Persistence in different decades

An ideal use for the four measures of persistence barriers is investigating interdecadal variability in the NINO3 SST and SOI. Using a window of fixed length, we calculate the persistence map for a windowed segment of the dataset. We then change the starting point of the window in yearly increments. The four measures, A , τ , Δp , and p_b , then become a function of time and describe the interdecadal variability of the persistence barrier. The window should be just large enough to provide sufficient data for calculating the persistence. Too small of a window makes the fit sensitive to the small amount of data in the window and moving the window can change the best-fit measures drastically. On the other hand, one cannot describe changes on timescales smaller than the window size. We have found that a window of 20 yr works best for the ENSO data.

Figure 11 contains plots of each measure of the NINO3 SST applied to the persistence of a running 20-yr window of the data. The bars labeled error indicate the 68% confidence interval averaged over all plotted data points. The time axis of each plot, varying from 1906 to 1987, describes the center of the window. The dashed line indicates the value of each measure for the full time series. Due to the nonlinearity of the persistence itself, the measures from the full time series are not the same as the time average of the windowed measures.

The interdecadal variation in the NINO3 SST measures contains four regimes of behavior (Fig. 11). Dotted lines mark the edges of the four intervals. Because the measures are calculated from a running window, we have fixed these boundaries at convenient dates. They




FIG. 11. Interdecadal variation of NINO3 SST as exhibited by A , τ , Δp , and p_b , using a 20-yr running window. The bars labeled error indicate a typical value of the 68% confidence interval. The dashed line is the value of the measure for the full time series. The dotted lines mark the boundaries between regions of qualitatively different behavior. (a) The strength A , (b) the decorrelation time τ , (c) the barrier width Δp , and (d) the barrier position p_b .

can readily be moved a few years in either direction and the conclusions will remain the same.

The first interval of interest in the NINO3 SST, from 1906 to 1915, contains a transition period. The barrier strength A decreases from 0.4 to 0.2 during this time

period. The stochastic decorrelation time τ and width Δp experience similar decreases. Here τ drops from near 2 yr to 0.4 yr, while Δp ranges from 0.12 yr to near zero. In contrast, the barrier position p_b increases from 0.3 to 0.38 yr, or from mid-April to late May.




FIG. 12. Persistence map of the NINO3 SST for a 20-yr window centered at 1906. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.927$, compared to $R^2 = 0.966$ for the full time series.

The next regime of interdecadal behavior is from 1916 to 1945. The strength and decorrelation time both remain low, with typical values of $A \approx 0.25$ and $\tau \approx 0.4$ yr. For most of the interval the width is around 0.1 yr, though it does later drop to near zero. The position remains in the late boreal spring, with p_b gradually decreasing from 0.38 to 0.3 yr.

One can see the difference between the 1906–15 and

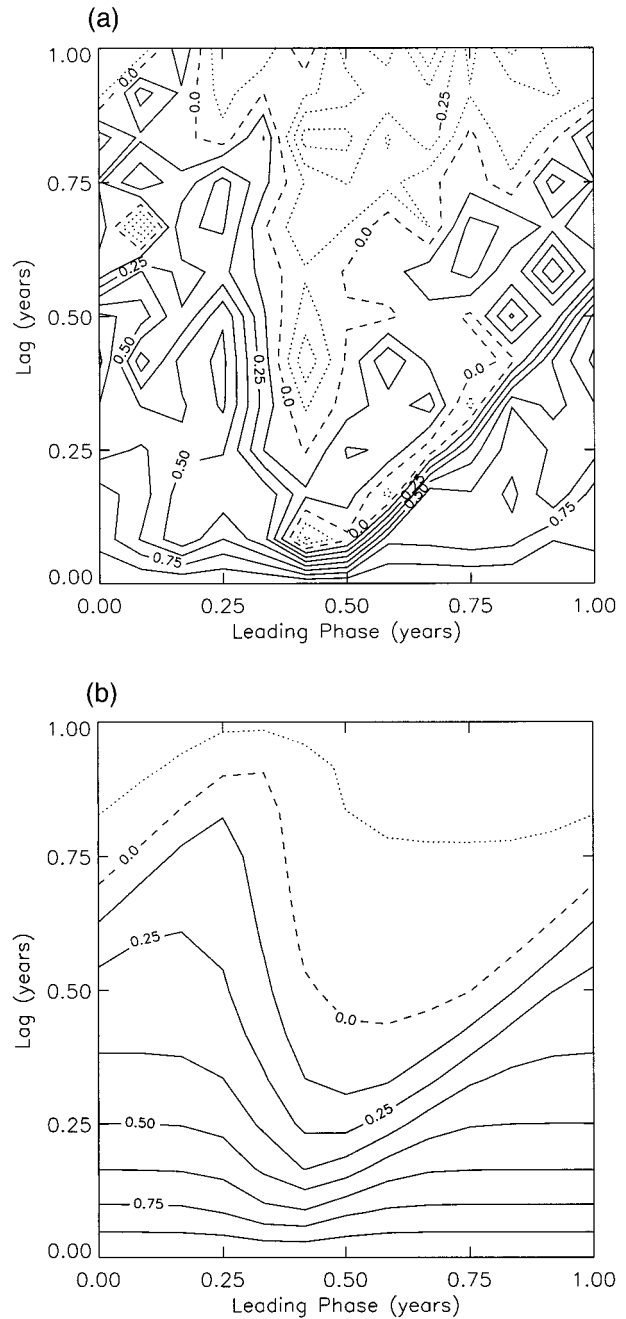


FIG. 13. Persistence map of the NINO3 SST for a 20-yr window centered at 1920. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.638$.

1916–45 intervals in Figs. 12 and 13. Figures 12a and 13a contain the persistence maps of the windowed data centered at 1906 and 1920, respectively. Figures 12b and 13b show the model [Eq. (10)] at the best-fit measures for the 1906 and 1920 maps, respectively. Also listed is the goodness of fit, R^2 , for each best-fit model. In addition to having a poorer goodness of fit, the 1920 map looks noisier than the 1906 map. The weakening

of barrier strength thus appears connected to the drop in stochastic decorrelation time as the noise in the barrier increases.

Between 1946 and 1975, there is a transition to a strong barrier, as exhibited by the large increase in A (Fig. 11). Here τ also increases, becoming large after 1960, when A reaches its largest value. Recall from section 4b that the stochastic term in the model barrier becomes constant at large τ , which is, in turn, connected to values of $A \geq 0.5$. The width remains near $\Delta p \approx 0.13$ yr, the value for the full NINO3 SST dataset. The position, however, drops below 0.2; that is, the barrier is now located in March.

Figure 14 contains an example of the 1946–75 period. The persistence map looks very much like the model barrier and has a large R^2 . This supports our interpretation of a strong persistence barrier during the 1946–75 interval.

During the final regime of interdecadal variation in the NINO3 SST measures, the barrier weakens. There is a drop in strength during 1976–87 to $A \approx 0.5$. The decorrelation time varies widely but is always larger than $\tau = 2$ yr. Although the large A and τ indicate a strong barrier, the width increases throughout the interval, from $\Delta p \approx 0.13$ yr to nearly triple that value by 1987. Last, the barrier position shifts back to May, with typically $p_b > 0.3$ yr.

Figures 15 and 16 exemplify the weakening persistence barrier in the NINO3 SST between 1976 and 1987. The persistence map from the window centered at 1980 (Fig. 15) shows a decrease in goodness of fit compared to the 1970 window (Fig. 14). The barrier for the 1987 window is wider still (Fig. 16) and also exhibits a slight decrease in R^2 relative to the 1980 window. The barrier is therefore currently very weak.

In summary, the interdecadal variations in the measures of the NINO3 SST persistence indicate four intervals of behavior in the barrier. From 1906 to 1915, the barrier is in decline, entering a period of weakness during 1916–45. In 1946–75, the persistence barrier grows strong. During the present interval, 1976–87, the barrier grows wide and weak.

Unlike the NINO3 SST, the windowed measures of the SOI contain no consistent intervals of behavior. Figure 17 shows the measures of the SOI for a 20-yr window, plotted in the same fashion as Fig. 11. The strength does increase after 1965, peaking at $A \sim 0.3$ near 1975, then decreasing down to 0.2 by 1987. Prior to 1965, however, A is simply noisy. The decorrelation time does have some consistent regimes of behavior, as seen in Fig. 17b. However, there is no overlap of these intervals with intervals of behavior in the other measures. For example, τ is at its lowest during 1925–45. However, the width Δp is near zero except between 1920 and 1935 and from 1955 to 1987.

Though we cannot draw many conclusions about barrier strength from Fig. 17, we do see interesting behavior on the barrier position. Prior to 1935, the SOI barrier




FIG. 14. Persistence map of the NINO3 SST for a 20-yr window centered at 1970. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.975$.

has a position sometime in May ($p_b \approx 0.35$ yr). It then shifts to $p_b \sim 0.1$ yr, or mid-February. After 1960, the barrier has a position in April, with $p_b \approx 0.25$ yr.

Part of the difficulty of interpreting the interdecadal variability of the SOI measures comes from the strong presence of noise in the persistence maps. One sees this in the persistence maps for the SOI windows centered at 1931 and 1933, Figs. 18a and 19a, respectively. Both




FIG. 15. Persistence map of the NINO3 SST for a 20-yr window centered at 1980. (a) Persistence map of the data; (b) The best-fit model barrier. The goodness of fit is $R^2 = 0.944$.




FIG. 16. Persistence map of the NINO3 SST for a 20-yr window centered at 1987. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.930$.

of these maps are similar to each other, as should be the case since the two windows differ only by four years of data. The best-fit models, however, differ drastically (Figs. 18b and 19b). For both data windows, the goodness of fit is poor, with R^2 well below its value for the full SOI. Thus, moving the data window in the SOI can cause drastic changes in the best-fit measures, making interpretation of the measures difficult.

Not all of the SOI data windows produce noisy bar-

riers that fit poorly to the model. The window centered at 1975, shown in Fig. 20, contains a strong barrier. Though still somewhat noisy, the persistence map visually resembles the best-fit barrier and the goodness of fit is near its value for the full time series. Note, too, that A reaches a maximum at 1975, with a decorrelation time of 0.6 yr. We therefore interpret the increase in A from 1965 to 1975 as a strengthening of the SOI persistence barrier.




FIG. 17. Interdecadal variation of the SOI as exhibited by A , τ , Δp , and p_b , using a 20-yr running window. The data is graphed in the same manner as in Fig. 11. (a) The strength A , (b) the decorrelation time τ , (c) the barrier width Δp , and (d) the barrier position p_b .

Contrasting Fig. 20, Fig. 21 shows the persistence map for the SOI window centered at 1987. The goodness of fit, R^2 , exhibits a decrease relative to the 1975 window, though it is not as low as R^2 for the 1931 and 1933 windows. The decrease in A between 1976 and 1987 thus heralds a weakening of the SOI persistence barrier. This coincides with the present weakening of the NINO3 SST barrier.

Several other works have examined the interdecadal variability of ENSO, with particular focus on the regime transition in the 1970s (Fig. 11). Balmaseda et al. (1995) found a strong spring decay in SST prediction skill in the 1970s, and a very weak seasonal dependence in prediction skill in the 1980s. They attributed this difference to strong phase locking of ENSO to the seasonal cycle in the 1970s, and weaker phase locking thereafter.





FIG. 18. Persistence map of the SOI for a 20-yr window centered at 1931. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.748$, compared to $R^2 = 0.841$ for the full time series.

A transition in the characteristics of the onset of El Niño in the late 1970s was noted by Wang (1995) and was attributed to an interdecadal change in the basic state of the SST. In their study of NINO3 SST, Torrence and Webster (1998) also found a transition in the late 1970s, with a strong persistence barrier before 1977 and a weaker one after.

Torrence and Webster (1998) also found transitions


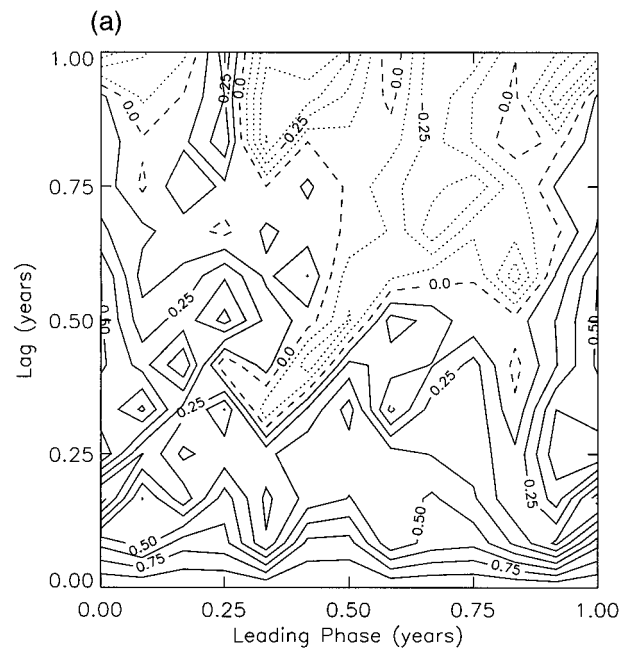


FIG. 19. Persistence map of the SOI for a 20-yr window centered at 1933. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.684$.

similar to the 1915 and 1945 transitions found here (Fig. 11). Balmaseda et al. (1995) noted several transitions throughout the early part of the century, some of which correspond to the transitions found here.

It is interesting to note that in their study of changes in the period of the Southern Oscillation, Wang and Wang (1996) find no significant transition in the 1970s. This is consistent with our finding that interdecadal var-




FIG. 20. Persistence map of the SOI for a 20-yr window centered at 1975. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.846$.

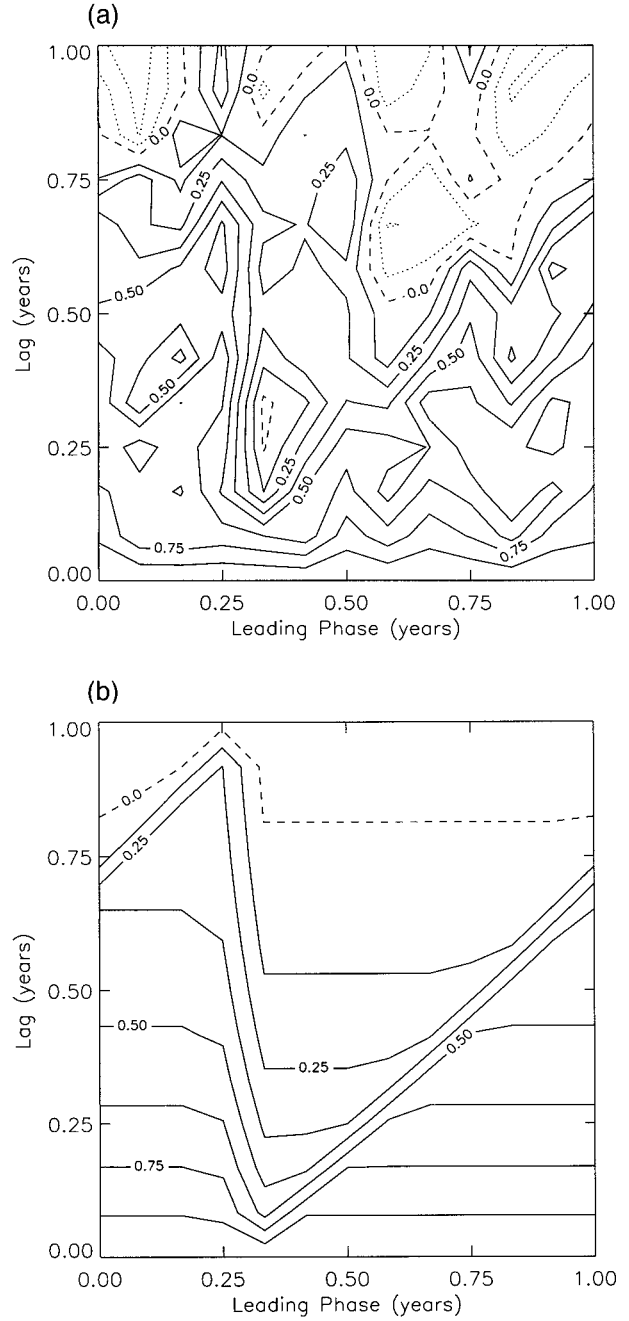


FIG. 21. Persistence map of the SOI for a 20-yr window centered at 1987. (a) Persistence map of the data; (b) the best-fit model barrier. The goodness of fit is $R^2 = 0.796$.

iability in the NINO3 SST barrier does not necessarily correspond with that in the SOI. Wang and Wang (1996) do, however, find two significant changes, one in the early 1910s in which the period lengthens from 3–4 yr to 5–7 yr, and one in the mid-1960s, in which the period falls to about 5 yr. These transitions do not seem to correspond to any particular feature in the interdecadal variability of the barrier in the SOI (Fig. 17), but they

do correspond roughly to the midpoints of the two regimes in the NINO3 SST where the barrier amplitude is changing significantly.

6. Summary

In this paper, we have analyzed the cyclostationary lag autocorrelation, or persistence, of a range of time

series with the goal of gaining insight into the causes and meaning of the persistence barrier seen in ENSO data. In some time series, the persistence is found to be a function of the phase of the year. We have focused our attention on one particular pattern of phase dependence, the pattern found in ENSO data, which we call a “persistence barrier.”

A time series composed of a single sine wave only produces a phase-dependent persistence map if the frequency of the wave is a biennial cycle or one of its harmonics, $\nu = (2n + 1)/2$. The phase dependence takes the form of a single discontinuous barrier if the wave is a biennial cycle, that is, $n = 0$. Otherwise, there are multiple discontinuous barriers in a single year.

A time series composed of two sine waves produces phase-dependent persistence maps under several circumstances. As expected, the persistence map is phase dependent when one or both frequencies are harmonics of the biennial cycle. However, phase dependence also results when neither frequency is a biennial harmonic, provided the frequencies are m^\pm complementary; that is, their sum or difference is an integer, m . Unit-complementary frequencies, frequencies whose sum is 1, produce an ENSO-like barrier, suggesting that barrier-producing harmonics are either biennial cycles or unit-complementary pairs of harmonics.

Time series with narrow continuous peaks can produce barriers when the peaks are centered on barrier-producing harmonics. If the phase angles of the harmonics vary rapidly, cancellation results in a persistence depending only on lag. A barrier is produced if the phase angles vary coherently.

Thus, in the time series considered here, the existence of a barrier depends on barrier-producing harmonics having coherent phase angles. The phase angles also determine the barrier position. This suggests that an arbitrary time series will only have a barrier if its dominant harmonics are barrier producing and have coherent phases.

Based on these results, we constructed a “model persistence barrier” to quantitatively measure the barrier. The model is a combination of an idealized barrier and the persistence of a red-noise, Ornstein–Uhlenbeck process. The barrier is described by four parameters: the barrier position, strength, and width, and the stochastic decorrelation time. The barrier in any time series is measured by a nonlinear least squares fit of the time series’ persistence to the model.

The measures show that the persistence of both the SOI and NINO3 SST datasets contains barriers that are statistically significant compared with the barrier produced by red noise. The predictability barrier is therefore an intrinsic, statistically significant property of the ENSO data.

Selective filtering experiments performed on these datasets demonstrate that the unit-complementary modes are the primary cause of the barrier, the modes nearest

the biennial mode being most important, and that the intraannual modes weaken the barrier. Phase-scrambling Monte Carlo experiments demonstrate that barrier-generating modes must also have coherent phase angles to produce a barrier.

The importance of phase coherency in the formation of a persistence barrier is also demonstrated in other work (Xue et al. 1994; Torrence and Webster 1998). Xue et al. (1994) show that the barrier is related to low cyclostationary variance. One can only have low cyclostationary variance, however, if the dominant harmonics of a time series have coherent or slowly varying phases. Only then will their superposition approach zero in a periodic fashion. Additionally, Torrence and Webster (1998) demonstrate a connection between changes in the ENSO persistence barrier and changes in whether ENSO phase locks to the annual cycle. It is likely that this phase locking to the annual cycle generates the necessary coherency between the unit-complementary harmonics in the SOI and NINO3 SST.

This paper concludes with a study of interdecadal variability of ENSO. The NINO3 SST persistence shows a weak barrier from 1915 to 1945 and another decline in the barrier after 1975. In contrast, the 1960s and 1970s show a strong barrier. This is consistent with other investigations of ENSO interdecadal variability (Balmaseda et al. 1995; Wang and Wang 1996; Torrence and Webster 1998). Additionally, the barrier appears to have shifted location, ranging from late April to mid-May during the first half of the century, shifting to late February to mid-March during the 1950s, 1960s, and 1970s, and returning to May recently.

Overall, the barrier in the SOI is quite different than that in the NINO3 SST. The SOI barrier is weaker, is less statistically significant, has a shorter stochastic decorrelation time, and, unlike the NINO3 SST, has no clear-cut interdecadal regimes. This is explained by the fact that the atmospheric component of the tropical climate system is noisier than the ocean and has shorter correlations. Thus, this result provides further support for the common practice of treating the atmosphere as a rapidly fluctuating random forcing of the ocean.

ENSO is commonly thought to be driven by a mode with a roughly 2–7-yr period. The existence of such a strong broad interannual mode is necessary for a barrier, but as indicated above, the time series must have other features as well. Thus, the existence and interdecadal variability of the ENSO persistence barrier puts stronger restrictions on ENSO dynamics than the mere existence of an ENSO mode. However, a detailed understanding of the connection between barriers and the underlying dynamics has not been achieved and is left for future work. Further, the persistence barrier has often been called a predictability barrier; without a detailed understanding of how the error-growth characteristics of the dynamics are related to barrier prop-

erties, it is unclear whether barriers truly indicate limited prediction.

Acknowledgments. SOI and NINO3 SST data from 1950–96 were courtesy of Don Garrett, Climate Prediction Center, NOAA. Data from 1896–1949 were courtesy of Chris Torrence, PAOS, University of Colorado. Thanks to Sue Haupt, Cécile Penland, Chris Torrence, and Peter Webster for discussions. Computations were performed at the PAOS computing facility. This work was funded by NOAA Grant DOC NA56GP0230.

APPENDIX A

Persistence of a Sine Wave

We first compute the persistence of Eq. (4), a single sinusoidal periodic oscillation. An important identity needed to analytically calculate the persistence in this and the other appendices is

$$\langle e^{\pm i2\pi\nu t} \rangle(p) = \begin{cases} 0 & \nu \neq n, \\ e^{\pm i2\pi\nu p} & \nu = n, \end{cases} \quad (A1)$$

where n is an integer. Letting $\beta \equiv -\phi/2\pi\nu$, one can now compute the persistence:

$$C(p, l) = \begin{cases} \text{undefined,} & \nu = n; \\ \frac{\sin[2\pi\nu(p - \beta)] \sin[2\pi\nu(p - \beta - l)]}{|\sin[2\pi\nu(p - \beta)] \sin[2\pi\nu(p - \beta - l)]|}, & \nu = \frac{2n + 1}{2}; \\ \cos(2\pi\nu l), & \nu \neq \frac{n}{2}. \end{cases} \quad (A2)$$

$$\nu = \frac{2n + 1}{2}; \quad (A3)$$

$$\nu \neq \frac{n}{2}. \quad (A4)$$

Note that in the second case, where ν is a half-integer, $C(p, l)$ can be written as

$$C(p, l) = \frac{x(p)x(p - l)}{|x(p)x(p - l)|}, \quad (A5)$$

a quantity divided by its absolute value. It is thus either +1 or -1, indicating perfect correlation or anticorrelation. This can, in turn, be rewritten in the form of Eq. (5):

$$C(p, l) = \Theta(p - \beta)\Theta(p - \beta - l), \quad (A6)$$

where $\Theta(x) = (-1)^{\text{Int}(x)}$ and $\text{Int}(x)$ is the integer part of x .

APPENDIX B

Persistence of Two Sine Waves

For the persistence of a time series composed of two sine waves,

$$x(t) = a_1 \sin(2\pi\nu_1 t + \phi_1) + a_2 \sin(2\pi\nu_2 t + \phi_2), \quad (B1)$$

there are again multiple cases, depending on the value of ν . Several of the results are too complex to express on a single line. To facilitate the discussion of these

results, note that persistence can be expressed in the form

$$C(p, l) = \frac{V(p, l)}{[V(p, 0)V(p - l, 0)]^{1/2}}, \quad (B2)$$

where the function $V(p, l)$ is, apart from a constant factor, the cyclostationary covariance, $\langle x(t)x(t-l) \rangle(p)$. We shall therefore state only $V(p, l)$ in place of the full persistence.

If either ν_1 or ν_2 is an integer, then that Fourier mode of the time series is removed by computing the anomaly. The persistence is then determined from Eqs. (A2)–(A4) applied to the remaining mode. When neither frequency is an integer, the persistence can still be found in closed form. For the case where $\nu_1 = n/2$ and $\nu_2 \neq n/2$, $V(p, l)$ has the form

$$V(p, l) = \hat{C} + \alpha^2 \cos[2\pi m(p - \beta - l/2)], \quad (B3)$$

where $\alpha = a_1/a_2$ is the ratio of amplitudes, $\hat{C} = \alpha^2 \cos(2\pi\nu_1 l) + \cos(2\pi\nu_2 l)$, and $\beta = -\phi_1/\pi m$. When both frequencies are half-integers, one can write C in the same form as the single-frequency, biennial case [Eq. (A5)]. It cannot, however, be reduced to Eq. (5).

The final case, where ν_1 and ν_2 are neither integers nor half-integers, falls into two subcases:

$$V(p, l) = \begin{cases} \hat{C}, & \nu_1 \pm \nu_2 \neq m; \\ \hat{C} + 2\alpha \cos(\pi\nu l) \cos[2\pi m(p - \beta - l/2)], & \nu_1 \pm \nu_2 = m; \end{cases} \quad (B4)$$

$$\nu_1 \pm \nu_2 = m; \quad (B5)$$

where m is an integer, $\beta = -(\phi_1 \pm \phi_2)/2\pi m$, $\hat{\nu} = \nu_1 \mp \nu_2$, and α and \hat{C} are defined as in Eq. (B3). Note that one can demonstrate relationships between negative values of m and phase shifts of various types. We shall ignore those in this paper.

The values of α and $\hat{\nu}$ determine how strong the gradients are in the persistence map. Specifically, as α nears 1, and $\hat{\nu}$ nears 0, the gradients of the continuous features grow stronger, becoming infinite at $\alpha = 1$. As might be expected, if $\alpha \gg 1$ or $\alpha \ll 1$, $C(p, l)$ loses phase dependence. When the amplitudes are equal, that is, $\alpha = 1$, the persistence defined by Eq. (B5) simplifies to

$$C(p, l) = \cos(\pi\hat{\nu}l)\Theta[m(p - \beta)]\Theta[m(p - \beta - l)], \quad (\text{B6})$$

which is the barrier of a single biennial harmonic modified by a phase-independent factor that oscillates with lag. Here $\hat{\nu}$ controls the ‘‘slope’’ of the steps. The steps are still discontinuous regardless of $\hat{\nu}$ or m . When the two frequencies are equal, $\nu_1 = \nu_2 = m/2$, then $\hat{\nu} = 0$, and the biennial result, Eq. (5), is exactly recovered.

APPENDIX C

Persistence of Narrow Peaks

Here we examine the persistence of Eq. (7),

$$x(t) = \int_0^\infty d\nu a(\nu) \sin[2\pi\nu t + \phi(\nu)], \quad (\text{C1})$$

under certain simplifying assumptions. Let $a(\nu) = a_1\eta(\nu - \nu_1) + a_2\eta(\nu - \nu_2)$, where $\eta(x)$ is a symmetric, positive-definite function with its maximum at zero and a width determined by a parameter, Δ . If ν_1 and ν_2 are unit complementary, $\nu_1 + \nu_2 = 1$, there is now a continuum of unit-complementary harmonics contained in the two peaks. Similarly, in the case of a single peak at $\nu_1 = \nu_2 = 1/2$, there is a continuum of unit-complementary harmonics surrounding this one peak.

We assume that the peaks are narrow, $\Delta \ll 1$, and that their shape is Gaussian, $\eta(\nu) = e^{-\nu^2/2\Delta^2}/\sqrt{2\pi\Delta^2}$. The latter assumption merely simplifies the form of the final result; any narrow distribution will suffice. We also assume that $\nu_1 < \nu_2$, $\nu_1 > \Delta$, $\nu_2 < 1 - \Delta$, and $\nu_1 + \Delta < \nu_2 - \Delta$. In other words, the two peaks neither overlap each other nor extend beyond the frequencies 0 and 1. Using these assumptions, one can show that the cyclostationary mean of Eq. (7) is approximately zero.

Deriving the cyclostationary covariance, $\langle x(t)x(t - l) \rangle_y(p)$ is lengthy and not straightforward. We therefore wish to describe some of its more pertinent points. The covariance contains a product of two integrals. We commute these integrals with the cyclostationary average and perform the cyclostationary mean on the integrand first. Using Eq. (A1) and trigonometric identities, one can then reduce one of the two integrals to a discrete

sum. We then make use of the symmetry and narrowness of the peaks to approximate the integral of the sum to a single integral from $-\infty$ to ∞ . After some additional reductions, one can compute $V(p, l)$ [Eq. (B2)]. One finds

$$V(p, l) \cong \hat{C}e^{-2\pi^2l^2\Delta^2} - \alpha \int_{-\infty}^{\infty} d\nu e^{-(\nu/\Delta)^2} [\cos(p_1 + \phi_{12}) + \cos(p_2 + \phi_{21})], \quad (\text{C2})$$

where $\alpha = a_1/a_2$ is the ratio of amplitudes, $\hat{C} = \alpha^2 \cos(2\pi\nu_1l) + \cos(2\pi\nu_2l)$, $p_i = 2\pi[p - l(\nu_i - \nu)]$, and $\phi_{ij} = \phi(\nu_i + \nu) + \phi(\nu_j - \nu)$. This equation can also be used for the case of a single biennial peak, $\nu_1 = \nu_2 = 1/2$, by setting $a_1 = a_2$.

The persistence can be further simplified in two different limits. If the phase angle $\phi(\nu)$ is rapidly varying over the width of a peak, then $\cos(p_i + \phi_{ij})$ acts like a random variable and the integral in Eq. (C2) vanishes. The persistence then becomes

$$C(p, l) \cong e^{-2\pi^2l^2\Delta^2} \frac{\hat{C}}{\alpha^2 + 1}. \quad (\text{C3})$$

This is the persistence of the two-mode noncomplementary frequency case, Eq. (B4), multiplied by an exponential decay in lag, a decay whose rate is governed by the width of the peak, Δ .

One can also simplify the persistence in the limit of slowly varying phase differences. Assuming $\phi''(\nu_i)\Delta^2 \ll 1$, and $\phi'(\nu_1) - \phi'(\nu_2) \ll 1$, where prime denotes differentiation with respect to ν , we Taylor expand the cosine terms in Eq. (C2) around $\nu = 0$. If one keeps terms up to $O(\Delta^2)$, Eq. (C2) reduces further to

$$V(p, l) \cong e^{-2\pi^2l^2\Delta^2} [\hat{C} - 2\alpha \cos(\pi\hat{\nu}l) \{ \cos[2\pi(p - \beta - l/2)] - \epsilon_2 \sin[2\pi(p - \beta - l/2)] \}], \quad (\text{C4})$$

where $\hat{\nu} = \nu_1 - \nu_2$, $\epsilon_2 = [\phi''(\nu_1) + \phi''(\nu_2)]\Delta^2/4$, and $\beta = -[\phi(\nu_1) + \phi(\nu_2)]/2\pi$. This expression is similar to the two-mode unit-complementary case, Eq. (B5), and reduces to it in the limit $\Delta \rightarrow 0$.

In the case of a single biennial peak, $\nu_1 = \nu_2 = 1/2$ and $a_1 = a_2$, Eq. (C2) again simplifies via a Taylor expansion around $\nu = 0$. However, we must now keep terms up to $O(\Delta^4)$, resulting in

$$V(p, l) \cong e^{-2\pi^2l^2\Delta^2} \left\{ 2 \sin[\pi(p - \beta)] \sin[\pi(p - \beta - l)] - [\epsilon_2(1 - 2\pi^2l^2\Delta^2) + \epsilon_4] \times \sin[2\pi(p - \beta - l/2)] + \frac{3}{2}\epsilon_2^2 \cos[2\pi(p - \beta - l/2)] \right\}, \quad (\text{C5})$$

where $\epsilon_4 = [\phi''''(\nu_1) + \phi''''(\nu_2)]\Delta^4/32$. As in the unit-complementary case, Eq. (C5) bears a strong resem-

blance to the persistence from a single biennial mode, Eq. (A3), and reduces to (A3) when $\Delta = 0$.

Note that the assumptions used to obtain Eqs. (C4) and (C5) do not strictly require slowly varying phase. The sole requirement is that $\phi'(\nu_1) - \phi'(\nu_2) \ll 1$. While this is certainly satisfied by slow variations, $\phi'(\nu_i) \ll 1$, it is also satisfied by $O(1)$ slopes in the phase, provided the slopes are similar at the two peaks. Thus, Eqs. (C4) and (C5) are valid when the phases of the two spectral peaks vary coherently.

REFERENCES

- Balmaseda, M. A., M. K. Davey, and D. L. T. Anderson, 1995: Decadal and seasonal dependence of ENSO prediction skill. *J. Climate*, **8**, 2705–2715.
- Battisti, D. S., and A. C. Hirst, 1989: Interannual variability in a tropical atmosphere–ocean model: Influence of the basic state, ocean geometry and nonlinearity. *J. Atmos. Sci.*, **46**, 1687–1712.
- Bjerknes, J., 1966: A possible response of the atmospheric Hadley circulation to equatorial anomalies of ocean temperature. *Tellus*, **18**, 820–829.
- , 1969: Atmospheric teleconnections from the equatorial Pacific. *Mon. Wea. Rev.*, **97**, 163–172.
- Cane, M., and S. E. Zebiak, 1985: A theory for El Niño and the Southern Oscillation. *Science*, **228**, 1085–1086.
- Chen, D., S. E. Zebiak, A. J. Busalacchi, and M. A. Cane, 1995: An improved procedure for El Niño forecasting: Implications for predictability. *Science*, **269**, 1699–1702.
- Gardiner, C. W., 1990: *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. 2d ed. Springer-Verlag, 442 pp.
- Goldberg, D. E., 1989: *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 412 pp.
- Haupt, R. L., and S. E. Haupt, 1998: *Practical Genetic Algorithms*. John Wiley and Sons, 192 pp.
- Huang, J.-P., and G. R. North, 1996: Cyclic spectral analysis of fluctuations in a GCM simulation. *J. Atmos. Sci.*, **53**, 370–379.
- Kvålseth, T. O., 1985: Cautionary note about R^2 . *Amer. Stat.*, **39**, 279–285.
- Moré, J. J., B. S. Garbow, and K. E. Hillstom, 1980: User guide for MINPACK-1. Rep. ANL-80-74, Argonne National Laboratory, Argonne, IL, 261 pp. [Available from Argonne National Library, 9700 S. Cass Avenue, Argonne, IL 60439.]
- Oort, A. H., Y. H. Pan, R. W. Reynolds, and C. F. Ropelewski, 1987: Historical trend in the surface temperature over the oceans based on the COADS. *Climate Dyn.*, **2**, 29–38.
- Philander, S. G., 1990: *El Niño, La Niña, and the Southern Oscillation*. Academic Press, 293 pp.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, 1992: *Numerical Recipes in C: The Art of Scientific Computing*. 2d ed. Cambridge University Press, 994 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.
- Ryan, T. P., 1997: *Modern Regression Methods*. John Wiley and Sons, 515 pp.
- Torrence, C., and P. Webster, 1998: The annual cycle of persistence in the El Niño–Southern Oscillation. *Quart. J. Roy. Meteor. Soc.*, **124**, 1985–2004.
- Troup, A. J., 1965: The “southern oscillation.” *Quart. J. Roy. Meteor. Soc.*, **91**, 490–506.
- Vallis, G. K., 1988: Conceptual models of El Niño and the Southern Oscillation. *J. Geophys. Res.*, **93** (C), 13 979–13 991.
- Wang, B., 1995: Interdecadal changes in El Niño onset in the last four decades. *J. Climate*, **8**, 267–285.
- , and Y. Wang, 1996: Temporal structure of the Southern Oscillation as revealed by waveform and wavelet analysis. *J. Climate*, **9**, 1586–1598.
- Webster, P. J., 1995: The annual cycle and the predictability of the tropical coupled ocean–atmosphere system. *Meteor. Atmos. Phys.*, **56**, 33–55.
- , and S. Yang, 1992: Monsoon and ENSO: Selectively interactive systems. *Quart. J. Roy. Meteor. Soc.*, **118**, 877–926.
- Wright, P. B., 1979: Persistence of rainfall anomalies in the central Pacific. *Nature*, **277**, 371–374.
- , 1985: The Southern Oscillation: An ocean–atmosphere feedback system? *Bull. Amer. Meteor. Soc.*, **66**, 398–412.
- Xue, Y., M. A. Cane, S. E. Zebiak, and M. B. Blumenthal, 1994: On the prediction of ENSO: A study with a low-order Markov model. *Tellus*, **46A**, 512–528.
- Zebiak, S., and M. Cane, 1987: A model El Niño–Southern Oscillation. *Mon. Wea. Rev.*, **115**, 2262–2278.