A nonlinear time series model of El Niño

A. D. Hall[1], J. Skalin[2] and T. Teräsvirta[3]

September 1998

Stockholm School of Economics
Working Paper Series in Economics and Finance, No. 263

Abstract. A smooth transition autoregressive model is estimated for the Southern Oscillation Index, an index commonly used as a measure of El Niño events. Using standard measures there is no indication of nonstationarity in the index. A logistic smooth transition autoregressive model describes the most turbulent periods in the data (these correspond to El Niño events) better than a linear autoregressive model. The estimated nonlinear model passes a battery of diagnostic tests. A generalised impulse response function indicates local instability, but as deterministic extrapolation from the estimated model converges, the nonlinear model may still be useful for forecasting the El Niño Southern Oscillation a few months ahead.

Keywords: Smooth transition autoregression; Nonlinearity; Time series model; El Niño; Southern Oscillation.

JEL Classification Codes: C22

Acknowledgements: The research on this paper was initiated while Hall was a visitor at the Stockholm School of Economics and he wishes to acknowledge its hospitality and financial support. Skalin and Teräsvirta acknowledge financial support from the Bank of Sweden Tercentenary Foundation. The authors wish to thank Michael McAleer for his constructive comments on an earlier draft.

[1]Corresponding author.
School of Finance and Economics, University of Technology, Sydney
PO Box 123, Broadway NSW 2007, Australia.
Phone: +61 2 9281 2020, Fax: +61 2 9281 0364. Email: Tony.Hall@uts.edu.au

[2,3]Department of Economic Statistics, Stockholm School of Economics,
PO Box 6501, SE-113 83 Stockholm, Sweden
Email: Joakim.Skalin@hhs.se and Timo.Terasvirta@hhs.se

1. INTRODUCTION

El Niño Southern Oscillation (ENSO) is a disruption of the ocean-atmosphere system in the tropical Pacific ocean that has important consequences for global weather conditions. The common-day usage of the term El Niño refers to the extensive warming of the central and eastern Pacific Ocean. In normal non-El Niño conditions trade winds blow towards the west across the tropical Pacific piling up warm surface water in the west Pacific. As a result the sea surface is higher at Indonesia than at Ecuador, and the sea surface temperature (SST) is about 8 degrees Centigrade higher in the west. The cool temperatures off South America are due to an upwelling of cold water from deeper levels, which is nutrient-rich, supporting high levels of primary productivity, diverse marine ecosystems, and major fisheries. As rainfall is found in rising air over the warmest water, the east Pacific is relatively dry. During El Niño the trade winds weaken in the central and western Pacific leading to a depression of the thermocline in the eastern Pacific, cooling the surface and cutting off the supply of nutrient rich thermocline water. The result is a rise in sea surface temperature and a drastic decline in primary productivity. Rainfall follows the warm water eastward, with associated flooding in Peru and drought in Indonesia and Australia. The eastward displacement of the atmospheric heat source overlaying the warmest water results in large changes in the global atmospheric circulation, which in turn causes changes in weather in regions far removed from the tropical Pacific. La Niña is the other phase, when sea surface temperatures in the central and eastern tropical Pacific are unusually low and when the trade winds are very intense. Together, these two natural processes form the ENSO, and on average, ENSO events occur every 4.5 years, with a range of 2 to 10 years.

The British scientist Sir Gilbert Walker while on assignment in India trying to find a way to predict the Asian monsoon, discovered a remarkable connection between barometer readings at stations on the eastern and western sides of the Pacific. When pressure rises in the east, it usually falls in the west, and vice versa. Walker coined the term Southern Oscillation to

dramatise the ups and downs in this east-west seesaw in Southern Pacific barometers. Walker noticed that “low-index” seasons, when pressure is low on the eastern side of the Pacific and high on the western side, are often marked by drought in Australia, Indonesia, India, and parts of Africa, and unusually mild winters in western Canada. In the late 1960s, University of California professor Jacob Bjerknes was the first to see a connection between unusually warm sea-surface temperatures and the weak easterlies and heavy rainfall that accompany low-index conditions. Ultimately, Bjerknes' discovery led to the recognition that the warm waters of El Niño and the pressure seesaw of Walker's Southern Oscillation are part of the same phenomenon. For more historical details see Gudmundson (1996), Philander (1990) and Allan, Lindesay and Parker (1990).

The use of the term El Niño has now changed and it is now primarily associated with ecological and economic disasters that coincide with torrential flooding in the eastern tropical Pacific, devastating droughts over the western tropical Pacific and unusual weather patterns over various parts of the world. The 1982-83 El Niño was by many measures the strongest thus far this century and The New York Times of 2 August 1983 catalogues detailed estimates of the worldwide economic impact of this occurrence. For instance, this El Niño has been blamed for between 1,300 and 2,000 deaths as well as over $US13 billion in damages to property and livelihoods. The effects in Australia, Africa and Indonesia included droughts, dust storms and bushfires. In Peru, areas where the normal seasonal rainfall was 6 inches had as much as 11 feet of rainfall causing devastating floods. As well, outbreaks of encephalitis on the East Coast of the USA, bubonic plague in New Mexico, increased rattlesnake bites in Montana, more shark attacks off the Oregon coast and the failure of the salmon harvest in the Eastern US have all been attributed to this El Niño event. There is even a suggestion that the angular momentum of the earth shifted slightly as a result of changes in the trade winds resulting in the length of the day increasing by 0.2 milliseconds. The weak El Niño of 1997-98 has been blamed for as diverse problems as the forest fires and resulting pollution in south

east Asia and for the luxuriant rough at the Olympic golf course in San Francisco, host of the 1998 US Open Golf Tournament.

There are two types of forecasts of ENSO events, dynamic forecasts based on numerical models of coupled ocean/atmospheric systems (for reviews see Allan, Lindesay and Parker, 1996, and Rothstein and Chen, 1996) and statistical forecasts based on historical records of sea surface temperature and barometric pressure differences. Records for SST in the Pacific ocean are very sparse prior to the decade long Tropical Ocean/Global Atmosphere experiment that began in 1985. On the other hand, barometric pressure has been recorded for over 100 years in many locations around the Pacific ocean. Each of these approaches have had some success in predicting ENSO events out to a lead time of about 1 year. However, it appears to be very difficult to accurately forecast outcomes past the boreal spring, indicating a less than complete understanding of the complex dynamic structure of the system (see Barnett et. al., 1988, Solow, 1995, Chen et. al., 1995 and Rothstein and Chen, 1996).

This paper will focus on modelling ENSO events using smooth transition autoregressive (STAR) time series methods. Shuzhen et al., (1988) use both open loop autoregressions and self exciting threshold autoregressions in modelling spatial SST averages, and Solow (1995) describes El Niño events in terms of a marked point process using Markov processes. Tziperman et al. (1994) suggest the ENSO cycle can be modelled as a low order chaotic process driven by the seasonal cycle. A casual inspection of the Southern Oscillation Index, a popular measure of ENSO events, indicates a possible assymetry in the sense that downturns in the data (that is occurrence of an El Niño event) appears to occur more rapidly that the recovery. STAR models are particularly suited to modelling this type of nonlinearity.

Section 2 reviews the modelling cycle for STAR models, section 3 describes the data and reports the results and section 4 concludes.

2. THE STAR MODEL AND TESTING CYCLE

2.1 The model

Our aim is to obtain a useful characterisation of the dynamics of the series. There are other ways to model nonlinearity, but the STAR family of models does have some useful properties. This model class is suitable for series with asymmetric cyclical variations and turbulent periods (see for instance Teräsvirta and Anderson (1992) and Teräsvirta (1995)), the estimated locally linear models have a simple interpretation, and a modelling cycle for the specification, estimation and evaluation stages already exists (see Teräsvirta, 1994, 1998).

Bacon and Watts (1971) were the first ones to define and apply a smooth transition model. In this paper, the STAR model is defined as

$$y_t = \pi_{10} + \pi_1' w_t + (\pi_{20} + \pi_2' w_t) F(y_{t-d}) + u_t \qquad (1)$$

where $\pi_j = (\pi_{j1},...,\pi_{jp})'$, $j = 1,2$, $w_t = (y_{t-1},...,y_{t-p})'$, *and* $u_t \sim NID(0,\sigma_u^2)$. The transition function $F(y_{t-d})$ is defined to be either a logistic function

$$F(y_{t-d}) = (1 + exp\{-\gamma_L (y_{t-d} - c_L)\})^{-1},\ \gamma_L > 0 \qquad (2)$$

or an exponential function

$$F(y_{t-d}) = 1 - \exp\{-\gamma_E (y_{t-d} - c_E)^2\},\ \gamma_E > 0 \qquad (3)$$

Model (1) with transition function (2) is a logistic STAR model of order *p*, LSTAR(*p*), whereas (1) with (3) is an exponential STAR model of order *p*, ESTAR(*p*). Other models are special cases of the STAR specifications. The LSTAR model approaches a two-regime

threshold autoregressive model (see Tong (1990)) when $\gamma_L \to \infty$, since (2) in the limit is a step function of y_{t-d} the value of which changes from zero to unity at c_L. When $\gamma_L \to 0$, the LSTAR model approaches a linear AR(*p*) model. On the other hand, the ESTAR model, equation (1) approaches a linear model both as $\gamma_E \to 0$ and (with probability one) as $\gamma_E \to \infty$. However, by a suitable modification one obtains a STAR model which approaches a three-regime threshold autoregressive model whose outer regimes are identical as $\gamma_E \to \infty$, see, for example, Teräsvirta (1998). If, $c_E = \pi_{20} = 0$, then the ESTAR model is identical to the exponential autoregressive model of Haggan and Ozaki (1981).

The role of the transition function in (1) is that it allows the coefficients for lagged values of y_t, $[\pi_1 + \pi_2 F(y_{t-d})]'$, and the intercept, $[\pi_{10} + \pi_{20} F(y_{t-d})]$, to change smoothly with y_{t-d}, so that the local dynamics of the model change with y_{t-d}. The LSTAR model allows the local dynamics to be different for high and low values of the transition variable, y_{t-d}. The modelling of local dynamics as a function of a lagged value of *y* makes it possible to model nonlinear effects of a shock. For instance, if a negative shock pushes a realization away from a locally stable regime (*F* close to unity, say), the subsequent change in the value of *F* changes the local dynamics (*F* now close to zero, say). If this regime contains a pair of explosive complex roots, *y* may be returned to the previous level a lot more quickly than would be the case if it followed a linear AR process. In contrast to the LSTAR case, the ESTAR transition function is symmetric about c_E in the sense that the local dynamics are the same for high as for low values of y_{t-d}, whereas the mid-range behaviour of the variable (values close to c_E) is different. The mid-regime does not necessarily have to be locally stable, because with the exponential transition function, it is possible for *y* to move rapidly between very small and very large values for which local dynamics are stable. In the modelling cycle, we allow the data decide which of the types of STAR models we fit to series

for which linearity is rejected. Diagnostic tests will reveal whether a STAR model offers an adequate characterization of the data or not.

2.2 Testing linearity and evaluating STAR models

The modelling cycle for building STAR models is discussed in Teräsvirta (1994, 1998). Testing linearity against STAR constitutes the first step of the model specification stage. In order to do that one first selects a linear autoregressive model for the series with apparently no autocorrelation in the residuals, by applying an appropriate model selection criterion such as the Akiake Information Criterion (AIC). The selected linear autoregressive model is the null model. For details of the test (F_L) with power against both LSTAR and ESTAR, see Teräsvirta (1994). The test is carried out for different values of the unknown delay parameter *d*, and the value of *d* associated with the test with the smallest *p*-value is selected. If none of the *p*-values is sufficiently small, linearity is not rejected. Note that if testing linearity were the main point of the whole investigation we could assume *d* unknown and carry out the test starting from that assumption, as in Luukkonen, Saikkonen and Teräsvirta (1988), and so control the overall significance level of the test. In this paper model selection, including the choice of *d*, is an important part of the work so we test linearity conditionally on *d* and also use the results to select the delay. The choice between the LSTAR and ESTAR specifications is based on a sequence of nested hypothesis tests (F_4, F_3 and F_2) as detailed in Teräsvirta (1994).

The estimated model is evaluated by a series of tests. As usual, the assumption of no error autocorrelation is tested using the Lagrange-Multiplier test (LM_{AR}) that Eitrheim and Teräsvirta (1996) derived for this purpose. Their paper also contains two other tests. One is for testing the hypothesis of no remaining nonlinearity. In this test the alternative hypothesis is that the data-generating process is an additive STAR model with two 'STAR components' instead of a single one as in (2) or (3), and the tests of no remaining nonlinearity are based on

a third-order Taylor expansion of the second transition function. Finally, the constancy of the parameter vectors $(\pi_{j0}, \pi_j')', j = 1,2$, is tested against the hypothesis that the parameters change smoothly over time. Three tests are carried out. The first one, LM_1, assumes that the parameters change monotonically over time, the second one, LM_2, that the change is symmetric with respect to an unknown point in time, and the third one, LM_3, that the change is possibly non-monotonic but not necessarily symmetric. All these tests are carried out by auxiliary regressions (see Eitrheim and Teräsvirta, 1996).

3. DATA DESCRIPTION

The Southern Oscillation Index (SOI) is calculated from the monthly or seasonal fluctuations in the air pressure difference between Tahiti and Darwin. There are a number of different definitions used to calculate the SOI, but we employ the definition used by the Australian Commonwealth Bureau of Meteorology, namely, the Troup SOI. This is the standardised anomaly of the Mean Sea Level Pressure (MSLP) difference between Tahiti and Darwin. It is calculated as follows:

$$SOI = 10 \times \frac{(P_{DIFF} - P_{DIFFAVE})}{SD(P_{DIFF})}$$

where P_{DIFF} = Tahiti MSLP - Darwin MSLP, $P_{DIFFAVE}$ = long term average of P_{DIFF} for the month in question, and SD(P_{DIFF}) = standard deviation of P_{DIFF} for the month in question. Chen (1982) recommended the use of this particular index after a systematic analysis of various single-station and two-station MLSP indices of the Southern Oscillation. The monthly data from January 1876 is available from the Bureau of Meteorology web site. For this study we use the monthly SOI data from January 1876 until May 1998, and this is graphed in Figure 1. The statistical properties of the SOI are summarised by the correlogram, spectrum and normalised histogram in Figure 2. There is no indication of nonstationarity in the data series.

Negative values of the SOI indicate El Niño episodes, the most recent events in 1994/95 and 1997/98, whereas positive values of the SOI are associated with La Niña episodes and the most recent strong La Niña was in 1988/89.

4. MODELLING THE SERIES AND INTERPRETING THE RESULTS

In this section we report results from the estimation of STAR models for the series, the evaluation of the estimated models, and consider the dynamic properties of our models. Using the AIC a lag length of 13 is chosen for the linear autoregressive model. A summary of the results of the tests for nonlinearity can be found in Table 1. The table contains the lag length, p, of the autoregressive model chosen by AIC, the p-value of the test corresponding to values of the delay parameter, d, from 1 to 5. Linearity is rejected against STAR models, particularly with delay lengths of 1 and 3. For these delay values, the table also contains the results of the model selection test sequence for choosing between ESTAR and LSTAR as detailed in Teräsvirta (1994) and LSTAR is the preferred model. A number of LSTAR models with a variety of delay variables were estimated. Many of these were unsatisfactory, in that they failed our set of diagnostic tests. In particular the LSTAR(13) model for $d = 3$ was difficult to interpret and ARCH tests strongly rejected the null of no ARCH. This rejection is not taken to mean that the true model is the specified nonlinear model with ARCH errors, rather it is interpreted as misspecification of the conditional mean. The results were better for the LSTAR(13) model with $d = 1$. We choose to discuss the model in which the intercept in the linear part of the model is constrained to be equal to that in the nonlinear part of the equation, as this imposes the sensible restriction that in normal times the estimated mean of the SOI will be zero. Other restrictions on the parameters were strongly rejected at conventional significance levels, perhaps not surprisingly with in excess of 1400 observations. The preferred LSTAR(13) model is:

$$
\begin{aligned}
y_t = {} & -108.479 + 1.880y_{t-1} + 1.073y_{t-2} + 3.175y_{t-3} - 2.313y_{t-4} - 1.431y_{t-5} + 2.860y_{t-6} \\
& \quad (2.06) \quad (2.78) \quad (1.77) \quad (1.45) \quad (-2.07) \quad (-0.82) \quad (1.46) \\
& + 1.402y_{t-7} - 1.098y_{t-8} - 0.434y_{t-9} - 0.641y_{t-10} + 1.424y_{t-11} - 3.698y_{t-12} + 2.731y_{t-13} \\
& \quad (1.27) \quad (-0.74) \quad (-0.61) \quad (-0.50) \quad (1.25) \quad (-1.94) \quad (1.44) \\
& + \big[108.479 - 1.392y_{t-1} - 0.953y_{t-2} - 3.133y_{t-3} + 2.379y_{t-4} + 1.477y_{t-5} - 2.800y_{t-6} \\
& \quad (2.06) \quad (-2.06) \quad (-1.57) \quad (-1.43) \quad (2.12) \quad (0.84) \quad (-1.42) \\
& - 1.419y_{t-7} + 1.081y_{t-8} + 0.500y_{t-9} + 0.570y_{t-10} - 1.431y_{t-11} + 3.699y_{t-12} - 2.795y_{t-13} \big] \\
& \quad (-1.29) \quad (0.73) \quad (0.71) \quad (0.44) \quad (-1.25) \quad (1.94) \quad (-1.48) \\
& \times \big[1 + \exp\{ -9.858 \ (y_{t-1} + 27.294) / s(y) \} \ \big]^{-1} \\
& \qquad\qquad (-2.39) \quad (28.71)
\end{aligned}
$$

T = 1455, R^2 = 0.467, AIC = 4.120, s = 7.769, $s_{\text{nonlinear}}/s_{\text{linear}}$ = 0.993, sk = -0.071, ek = 0.974, LJB = 58.788 (p = 0.000), LM_{AR} =2.909 (p=0.021), LM_{ARCH} = 2.544 (p = 0.038)

The standardisation of the exponent of F by division by $s(y)$, the sample standard deviation of y_t, is introduced to make γ scale-free and thus facilitate the interpretation of its estimate. The estimated STAR model equation is reported together with a number of statistics: T is the sample size; AIC is the Akaike information criterion; s is the estimated standard deviation of the residuals; the ratio $s_{nonlinear} / s_{linear}$, where s_{linear} is the estimated standard error of the residuals from the linear autoregressive model used as a basis for linearity testing, which gives an idea of the relative gain in the fit from applying a LSTAR model instead of a linear autoregressive model; sk is the skewness, ek is the excess kurtosis and LJB is the test of normality suggested by Lomnicki (1961) and Jarque and Bera (1980); LM_{AR} is an Lagrange multiplier statistic of no autocorrelation based on four lags; and, LM_{ARCH} is an Lagrange multiplier statistic of no ARCH (Engle, 1982) based on four lags. The numbers in parentheses following values of test statistics are p-values, whereas those below the coefficient estimates are asymptotic "t-statistics" of the estimates.

The ratio $s_{\text{nonlinear}}/s_{\text{linear}}$ is quite close to unity which indicates that nonlinearity is only needed to characterise exceptional periods in the series and that the process thus is not inherently nonlinear, but the nonlinear model describes the most turbulent periods in the data better than the linear autoregressive model. The test of no error autocorrelation and no ARCH do not indicate misspecification. The autocorrelation function, spectrum and standardised histogram

of the residuals from the fitted LSTAR model are graphed in Figure 3. For the tests of no remaining nonlinearity, the p-values corresponding to each delay value of d are (d=1, p=0.339; d=2, p=0.089; d=3, p=0.030; d=4, p=0.170; and d =5, p=0.206), and given the number of tests, the smallest p-value of 0.030 is not very strong evidence against the model. The parameter constancy tests are reported in Table 2 and one of these tests provides some evidence of nonconstancy but this has not been followed up.

Figure 4 shows the shape of the transition function. Every point indicates an observation so that one can readily see which values the transition function has obtained and how frequently. The same information ordered over time is presented in Figure 5. Together these two figures indicate that the transition function normally has been close to unity. Also, as mentioned above, the estimated equation contains the parameter restriction $\pi_{10} = -\pi_{20}$, so that under the "normal regime" (F = 1) the local mean of the process equals zero. The intercept is only contributing to the process when F < 1, which is a rare event.

The dynamic behaviour of the model is characterised in two ways. Trying to interpret individual parameter estimates or the delay d does not give much useful information ($\hat{\gamma}$ is the only exception). It is more instructive to compute the roots of the characteristic polynomial at given values of the transition function F as in Teräsvirta (1994). The extreme values F = 0 and F = 1 are particularly interesting. The roots of the characteristic polynomial of given F = 0 and F = 1, respectively, can be found in Table 3. For F = 1, all roots are stationary, but for F = 0, there exists a real root which is greater than unity, and a number of unstable complex roots.

The stability of the estimated LSTAR model may also be investigated by using generalised impulse response functions. By doing so, the problems associated with traditional impulse response analysis applied to nonlinear models are avoided through averaging over 'histories', shocks, and 'futures'. The method used is that proposed by Koop, Pesaran and Potter (1996)

and we refer to that article for a general discussion. Here, we briefly describe the implementation of the method in this particular case. A random sample of 364 observations from the time series (approximately 25% of all available observations) is used as the set of 'histories'. For each history, 100 initial shocks are drawn randomly with replacement from – in this particular case – the full set of all residuals from the estimated LSTAR model. For each combination of history and initial shock, 800 replicates are generated, and the maximum forecast horizon is 10. That is, 800 prediction sequences (0, 1, … , 10 periods ahead) are computed according to the estimated model, both with the initial shock and without it using a residual, randomly drawn with replacement from all STAR residuals for each of the 800 replicates, in its place. The disturbance terms used in these predictions (the 'futures') are drawn randomly from the STAR residuals. For every one of the 11 horizons, the means over the 800 replicates are computed for the two prediction sequences (i.e. the sequences with and without the designated initial shock). The vector of differences between these two means constitutes an observation of the generalised impulse response. In this case, 364 histories and 100 initial shocks are used, so that 36400 pairs of history and initial shocks and 36400 generalised impulse response vectors of length 11 are generated.

In order to visualise the generalised impulse responses in a way that highlights their key distributional features, highest density regions are used as in Skalin and Teräsvirta (1998). (See Hyndman (1995, 1996) for a description of highest density regions and the density quantile method for estimating them). For each one of the 11 horizons, the density for that horizon is estimated with a kernel algorithm. The kernel routine used is based on Gauss code by King (1996) but has been modified to use a random sample from the computed impulse response values as the kernel estimation points instead of using equally spaced points. The 50% and 95% highest density regions are then estimated by applying the density quantile method to the kernel estimates. Figure 6 shows that the estimated stochastic LSTAR model is not stable. Initially the densities almost converge to a point as one would expect, but after six months start to become flatter and flatter over time. This behaviour suggests that the model

may only be used for short-term forecasting. The medians of the estimated forecast distributions would probably work well as point forecasts even at somewhat longer horizons, whereas the density-based tolerance intervals would not make good interval predictions. Note, however, that (biased) forecasts obtained by extrapolating the estimated LSTAR without noise do not explode, and these forecasts from June 1998 are graphed in Figure 7. One may conclude that the rare periods for which a nonlinear description is called for are really quite exceptional because the estimated LSTAR model accommodating the behaviour of the index during them displays explosive behaviour.

5. CONCLUSIONS

We can conclude that the estimated nonlinear LSTAR model describes the most turbulent periods in the data better than the linear autoregressive model. According to the model the dynamic behaviour of the ENSO during the turbulence is very different from that during "normal" times. On the other hand it is clear that because of the lack of stability the model has limited usefulness for forecasting purposes although it could still be used for predicting a few months ahead. Also note that probably no *univariate* model, ours included, can correctly predict an outbreak of El Niño, because the initial shock or shocks triggering such an event are exogenous to the univariate system. More information in the form of other, predictable, variables would be needed to improve this performance. Once the El Niño event is under way, however, the model may help predict the strength of the perturbance a few months ahead. Thus the main gain from this modelling exercise seems to be the improved understanding of the nonlinear dynamic behaviour of the ENSO.

References

Allan, R., Lindesay J. and Parker, D. (1996) *El Niño, Southern Oscillation and Climatic Variability*, CSIRO Publishing, Melbourne.

Bacon, D. W. and Watts, D. G. (1971) Estimating the transition between two intersecting straight lines, *Biometrika*, **58**, 525-534.

Barnett, T., Graham, N., Cane, M., Zerbiak, S., Dolan, S., O'Brien, J., and Legler, D. (1988) On the prediction of the El Niño of 1986-87, *Science*, **241**, 192-196.

Chen, D., Zebaic, S. E., Busalachhi, A. J. and Cane, M. A. (1995) An improved procedure for El Niño forecasting: implications for predictability, *Science*, **269**, 1699-1702.

Chen, W. Y. (1982) Assessment of Southern Oscillation sea-level pressure indices, *Monthly Weather Review*, **110**, 800-807.

Eitrheim, Ø. and Teräsvirta, T. (1996), Testing the adequacy of smooth transition autoregressive models. *Journal of Econometrics*, **74**, 59-75.

Engle, R. F. (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, **50**, 987-1007.

Haggan, V. and Ozaki, T. (1981) Modelling nonlinear random vibrations using an amplitude-dependent autoregressive time series model. *Biometrika*, **68**, 189-196.

Hyndman, R. J. (1995) Highest-density forecast regions for non-linear and non-normal time series models, *Journal of Forecasting*, **14**, 431-441.

Hyndman, R. J. (1996) Computing and graphing highest density regions, *The American Statistician*, **50**, 120-126.

Jarque, C. M. and Bera, A. K. (1980) Efficient tests for normality, homoscedasticity and serial independence of regression residuals, *Economics Letters*, **6**, 255-259.

Koop, G., Pesaran, M. H. and Potter, S. M. (1996) Impulse response analysis in nonlinear multivariate models, *Journal of Econometrics*, **74**, 119-147.

Lomnicki, Z. A. (!961) Tests for departure from normality in the case of linear stochastic processes, *Metrika*, **4**, 37-62.

Luukkonen, R., Saikkonen, P. and Teräsvirta, T. (1988) Testing linearity against smooth transition autoregressive models. *Biometrika*, **75**, 491-499.

Philander, S. G. (1990) *El Niño, La Niña and the Southern Oscillation*, Academic Press, London.

Rothstein, L. M. and Chen, D. (1996) The El Niño/Southern Oscillation phenomenon: seeking its "trigger" and working towards prediction, *Oceanus*, **39**, 39-41.

Shuzhen, P., Huiling, Y., Toole, J., Millard, B., McPhaden, M. J. and Mangum, L. J. (1990) Comparisons among autoregression models for forecasting El Niño events, *Air-Sea interaction in Tropical Western Pacific*, proceedings, US-PRC International TOGA Symposium, Beijing, 1988, China Ocean Press, 59-66.

Skalin, J. and Teräsvirta, T. (1998) Another look at Swedish business cycles, 1861-1988. *Journal of Applied Econometrics* (forthcoming).

Solow, A .R. (1995) An exploratory analysis of a record of El Niño events, 1800-1987, *Journal of the American Statistical Association*, **90**, 72-77.

Teräsvirta, T. (1994) Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89, 208-218.

Teräsvirta, T. (1995) Modelling nonlinearity in U.S. gross national product 1889-1987. *Empirical Economics*, **20**, 577-597.

Teräsvirta, T. (1998) Modeling economic relationships with smooth transition regressions. In *Handbook of Applied Economic Statistics*, eds Ullah, A. and Giles, D.E.A., 507-552, Dekker, New York.

Teräsvirta, T. and Anderson, H. (1992) Characterizing nonlinearities in business cycles using smooth transition autoregressive models. *Journal of Applied Econometrics*, **7**, S119-S136.

Tong, H. (1990) *Non-linear time series: a dynamical system approach*, Oxford University Press, Oxford.

Tziperman, E., Stone, L., Cane, M.A. and Jarosh, H. (1994) El Niño chaos: overlapping resonances between the seasonal cycle and the Pacific ocean- atmosphere oscillator, *Science*, **264**, 72-74.

Commonwealth Bureau of Meteorology, Australia home page, http://www.bom.gov.au/.

Gudmundson, C. (1996) El Niño and Climate Prediction, *Reports to the Nation*, University of Washington, http://atmos.washington.edu/gcg/rtn/.

King, G. (1996) "Dens.g, Gauss procedure for kernel density estimation", from URL http://gurukul.ucc.american.edu/econ/gaussres/.

Table 1: Results of linearity tests and model selection

Variable	*p*	*d*	$p(F_L)$	$p(F_4)$	$p(F_3)$	$p(F_2)$	STAR
SOI	13	1	0.033	0.076	0.315	0.074	L
SOI	13	2	0.108				
SOI	13	3	0.004	0.008	0.292	0.051	L
SOI	13	4	0.331				
SOI	13	5	0.058	0.106	0.443	0.071	L

Here *p* is the number of lags in the linear AR model, d is the delay factor, $p(F_L)$ is the p-value of the linearity test; $p(F_4)$, $p(F_3)$, $p(F_2)$ are the p-values of the tests in the model selection sequence, and the selected model family is (E = ESTAR, L = LSTAR). E is selected if $p(F_3)$ is the smallest value of the sequence , otherwise L is the choice (see Teräsvirta, 1994).

Table 2: Tests of Parameter Constancy for the LSTAR(13) Model

	All parameters	All linear parameters	All nonlinear parameters	Restricted intercept
LM3	1.187 (0.129)	1.212 (0.175)	1.130 (0.264)	0.418 (0.741)
LM2	1.232 (0.123)	0.993 (0.475)	0.944 (0.550)	0.567 (0.567)
LM1	1.657 (0.019)	1.504 (0.109)	1.441 (0.126)	0.868 (0.352)

p-values in parentheses

Table 3: Roots of characteristic polynomials for values of the transition function F.

Regime F = 0			Regime F = 1		
Root	Modulus (half-life)	Period	Root	Modulus (half-life)	Period
2.62	2.62		0.90 ± 0.13	0.91 (7.97)	42.4
-0.56 ± 1.24	1.36	3.15	-0.45 ± 0.71	0.84 (4.90)	2.94
-1.01 ± 0.31	1.05	2.21	0.61 ± 0.53	0.81 (4.32)	8.78
-0.57 ± 0.84	1.02	2.89	-0.73 ± 0.33	0.80 (4.06)	2.31
0.75 ± 0.67	1.01	8.62	0.30 ± 0.73	0.79 (3.86)	5.32
0.23 ± 0.82	0.85 (5.30)	4.82	-0.02 ± 0.77	0.77 (3.61)	3.93
0.80 ± 0.17	0.82 (4.42)	29.2	-0.73	0.73	

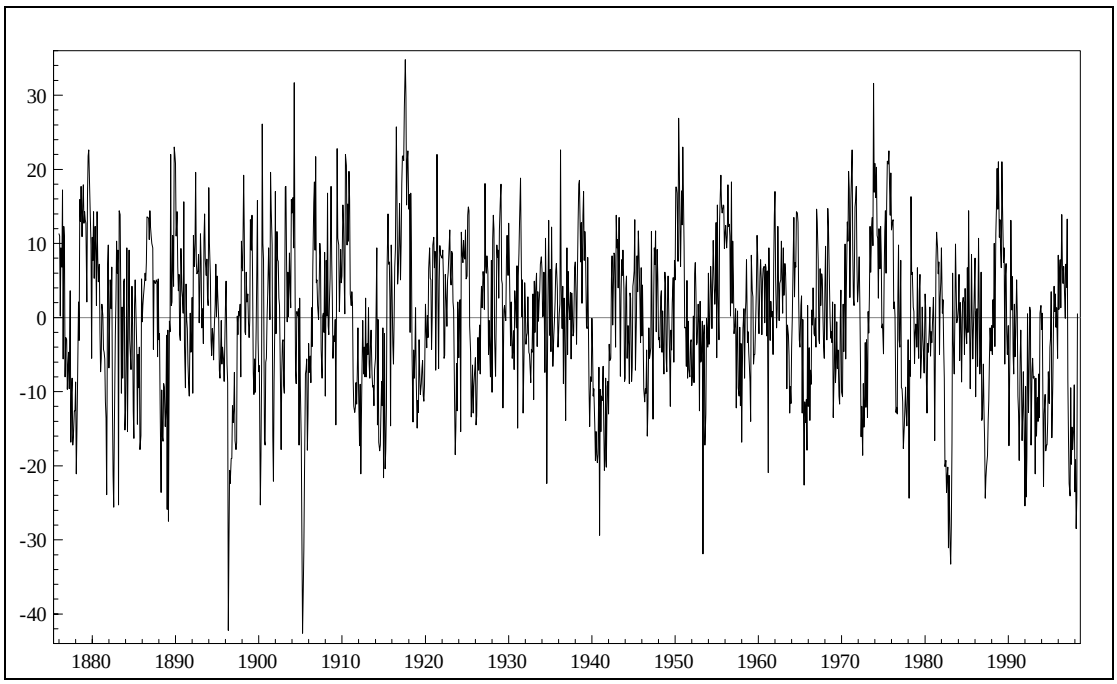

Figure 1: The Southern Oscillation Index (SOI)

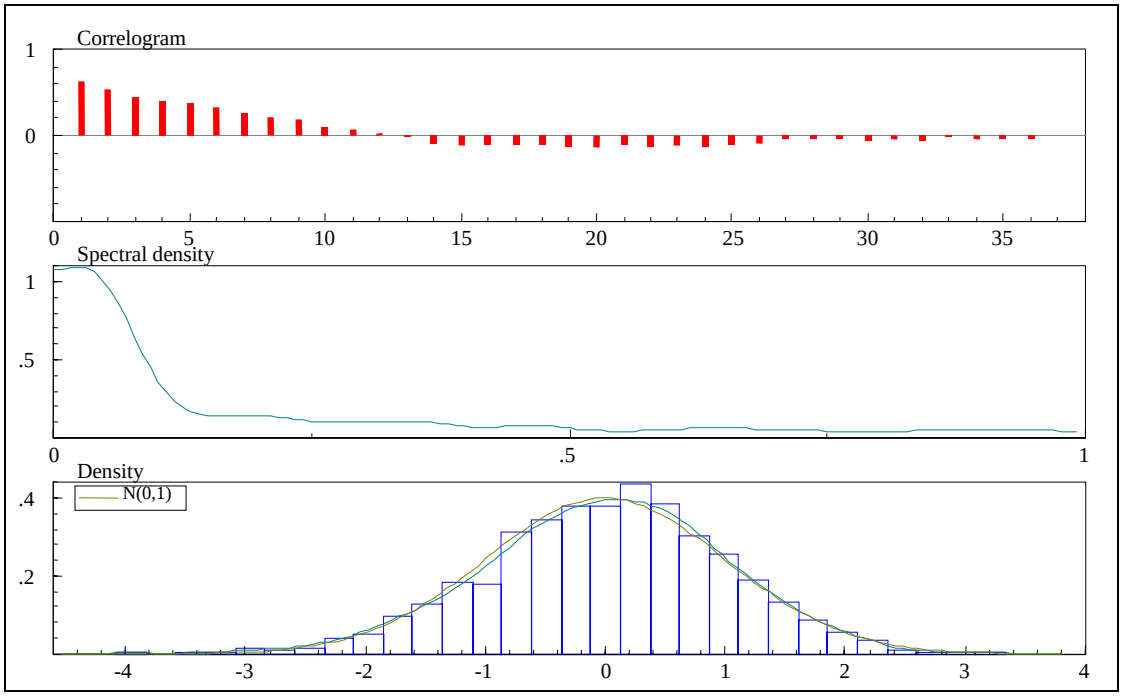

Figure 2: Summary Statistics For The SOI.

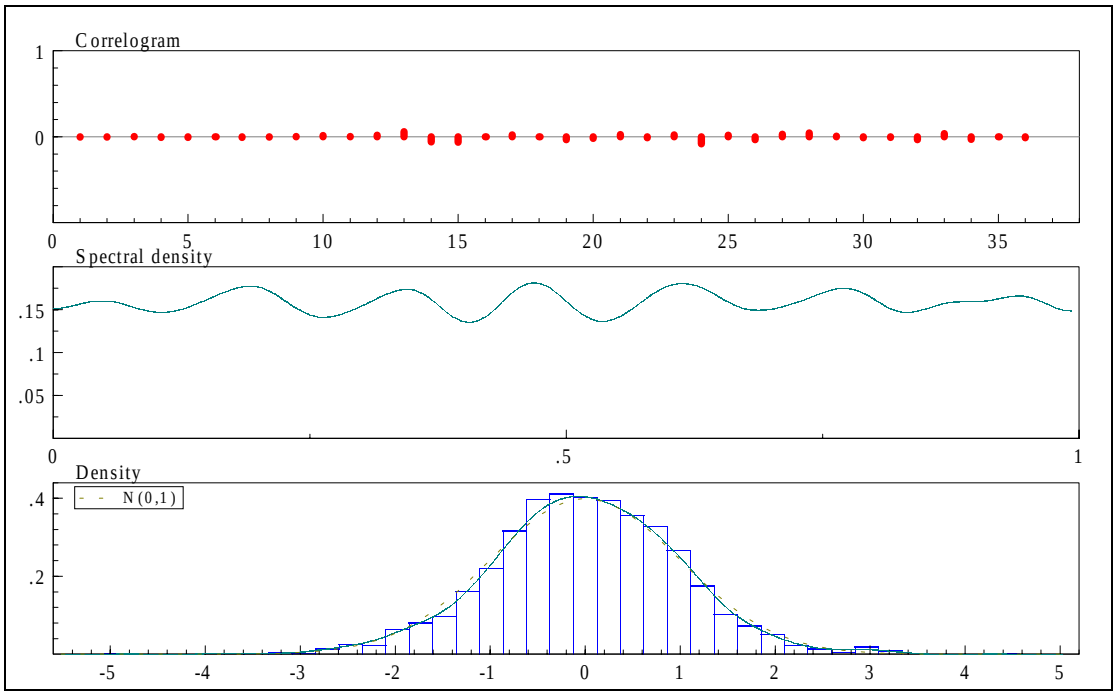

Figure 3: Summary Statistics For Residuals From The LSTAR(13) Model.

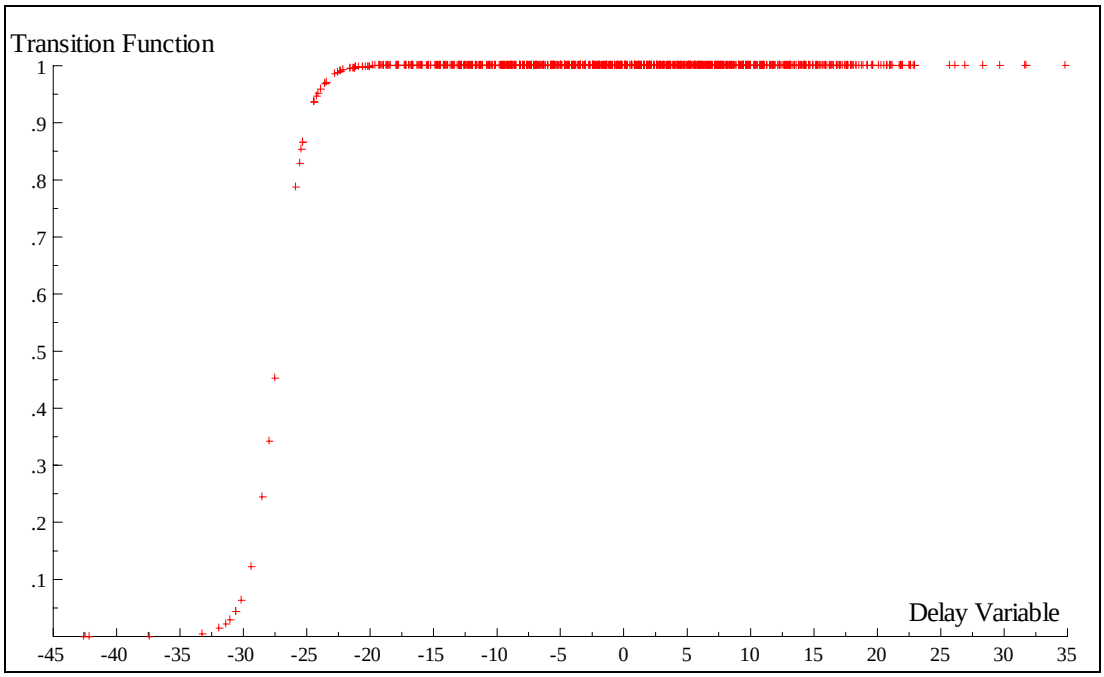

Figure 4: Transition Function As A Function Of The Observed Delay Variable Values

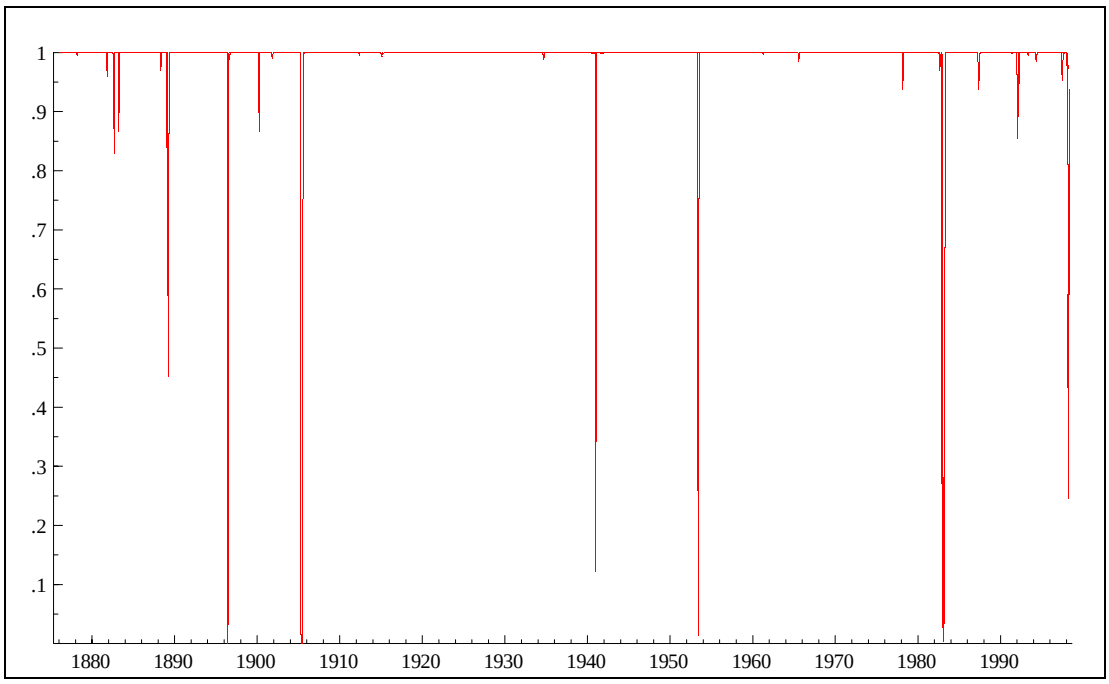

Figure 5: Transition Function From The LSTAR(13) Model.

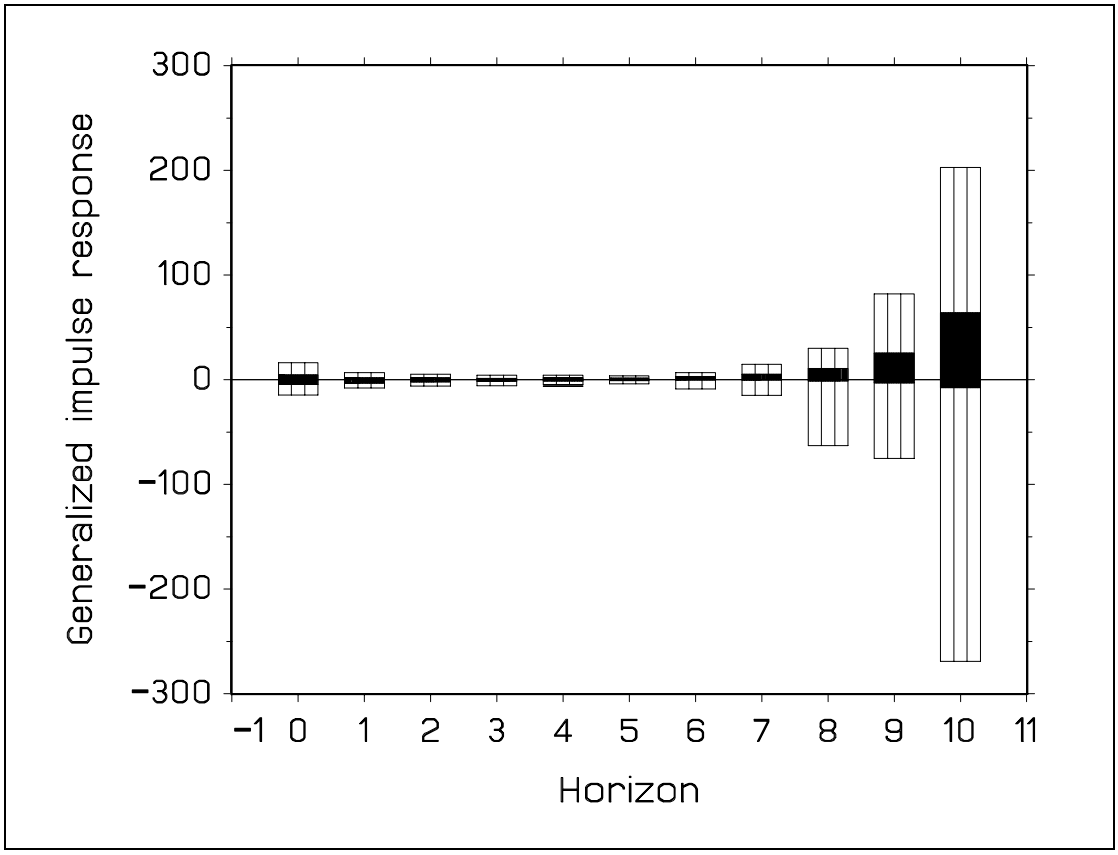

Figure 6: Generalised Impulse Response Functions (50% and 95% Higher Density Regions)

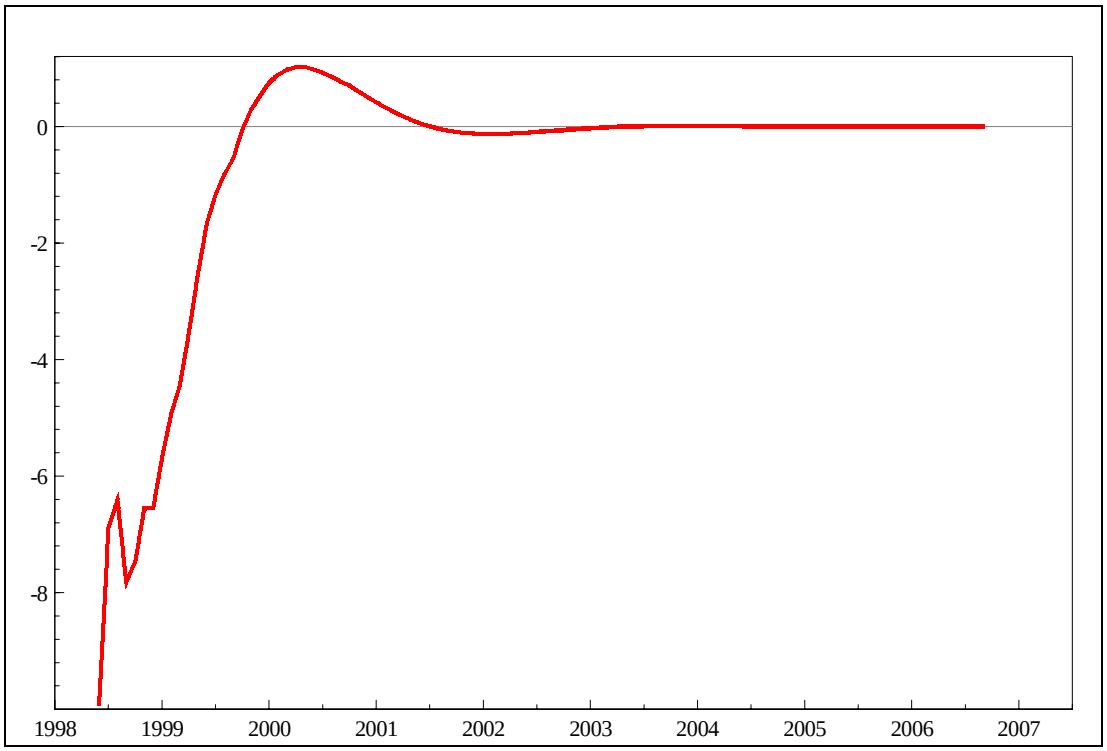

Figure 7: Deterministic Forecast Function from the LSTAR(13) Model