Institute of Actuaries of Australia

Quasi-Cyclic Risks

Prepared by Colin Priest

Presented to the Institute of Actuaries of Australia
XIV General Insurance Seminar 2003
9-12 November 2003

The Institute of Actuaries of Australia
Level 7 Challis House 4 Martin Place
Sydney NSW Australia 2000
Telephone: +61 2 9233 3466 Facsimile: +61 2 9233 3446
Email: insact@actuaries.asn.au Website: www.actuaries.asn.au

Quasi-Cyclic Risks

The Characteristics of Boom and Bust Processes

Colin Priest BEc FIAA
Colin_Priest@sigmaplus.com.au

Sigma Plus Consulting
www.sigmaplus.com.au

Introduction

At the 2001 General Insurance Seminar I presented a paper entitled “Regime Switching and Cycles” which used a Gibbs Sampling algorithm to fit parameters for a regime switching based saw-tooth model of building activity and the insurance cycle. During the question and answer period after the presentation, one of the questions from the floor asked how to determine whether a time series followed a quasi-cyclic pattern rather than a traditional time series process. There was no quick and simple answer to that question in that question and answer session.

This paper lists some of the characteristics of quasi-cyclic models. It then proposes some tests for whether the strength of certain quasi-cyclic characteristics are consistent with the well known autoregressive AR(1) and AR(n) processes. These tests are then applied to economic factors that are known to affect mortgage insurance claims.

In summary, the purpose of this paper is to show when to use autoregressive time series models, and when to consider using some of the lesser known quasi-cyclic processes.

Characteristics of Quasi-Cyclic Data Series

Not all quasi-cyclic data series share the same characteristics. A quasi-cyclic risk will have some of the following characteristics:

1. Consistent floor or ceiling values where the series turns around and heads the opposite direction
2. Long runs of consistently sloped movements in value, punctuated by sudden changes
3. A quasi-regular cycle length
4. Periods of stagnancy when values do not move
5. A possible unit root

A quasi-cyclic risk must *always* have the characteristic of a quasi-regular cycle length. This means that there is a concentration of the wavelengths within the observed time series. The wavelength is measured by the number of periods between successive peaks or troughs (points at which there is a change in slope). In a process with no memory, the distribution of wavelengths follows a Poisson distribution.

The real-life examples in this section of the paper show how some time series show different combinations of these characteristics.

Graph 1: Movement Between Limits – Seasonally Adjusted Building Approvals

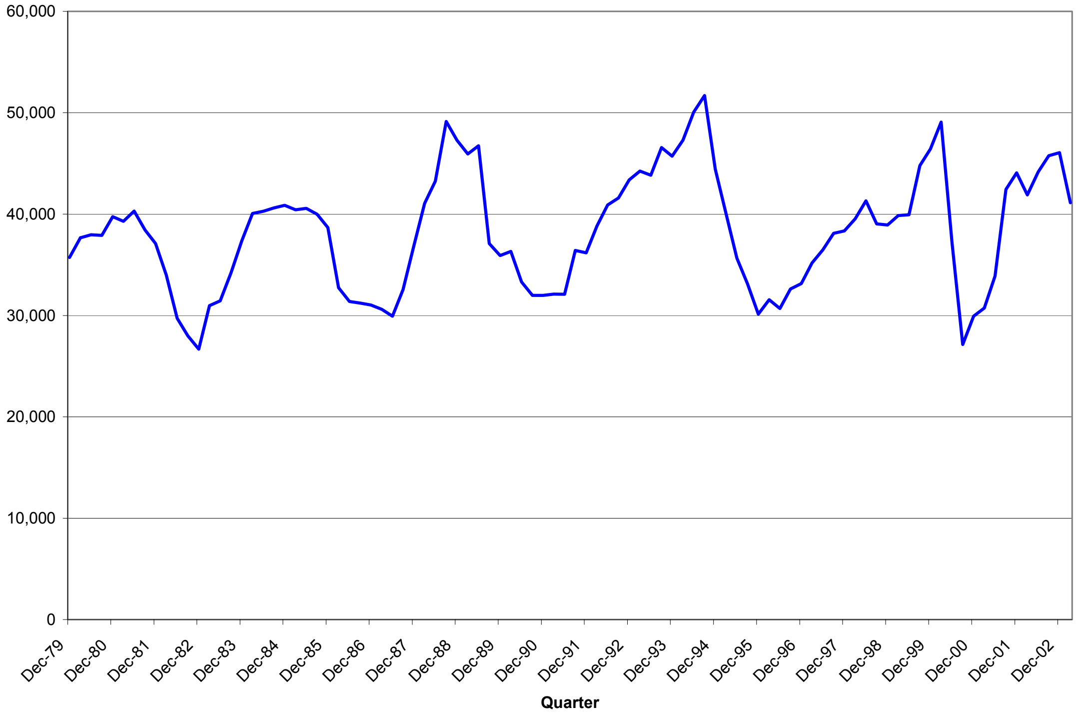

Many data series have a characteristic of a floor or ceiling for values, but this does not necessarily mean that they are quasi-cyclic. For example, unemployment rates cannot lie outside the range of 0% to 100% because the values are percentages of the population. With quasi-cyclic risks, we are more interested in series that move up and down between two limits. Australian building approvals counts are a clear example of this type of characteristic. Autoregressive time series processes do not have peaks and troughs with consistent heights, but instead have peaks and troughs over the entire range of possible data series values.

Graph 2: Periods of Consistent Slope – Slopes of Building Approvals

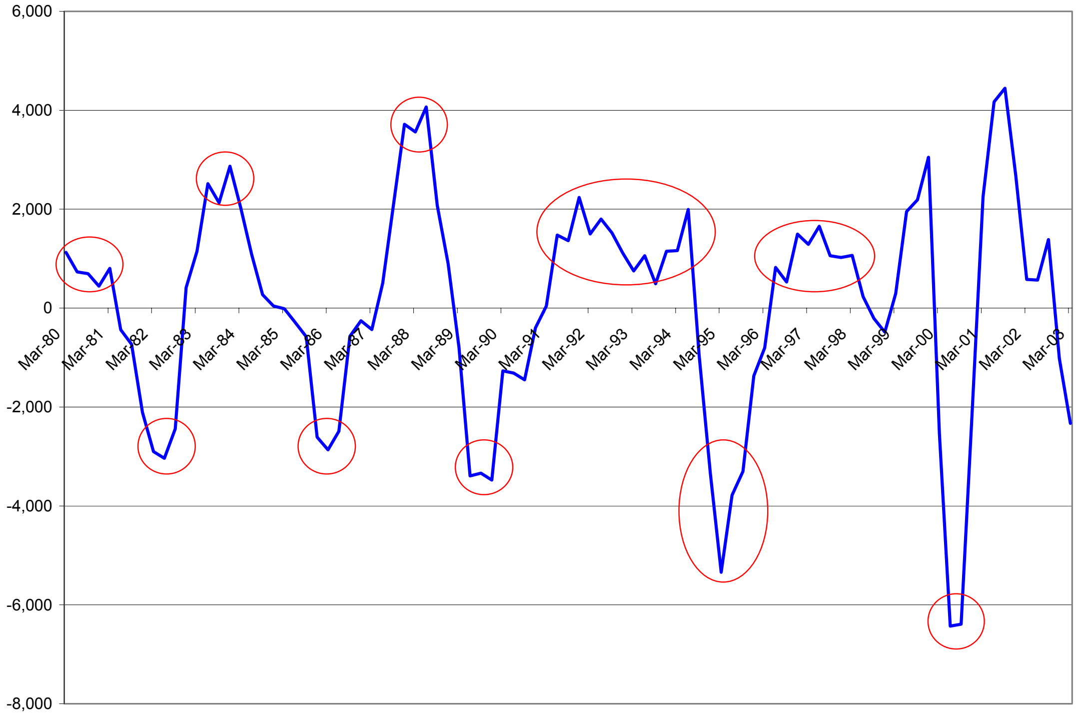

Autoregressive time series processes do not have long periods of consistently sloped movements in values. For example, in an AR(1) process the mean slope is related to the distance from the mean, and thus one would expect steeper slopes toward the mean when values have drifted away to extremes. Quasi-cyclic risks often have consistent changes in value over long periods. The slope of Australian Building Approvals shows a clear tendency towards long periods of upward or downward changes in value.

Graph 3: Quasi-Regular Cycle Lengths - Building Approvals

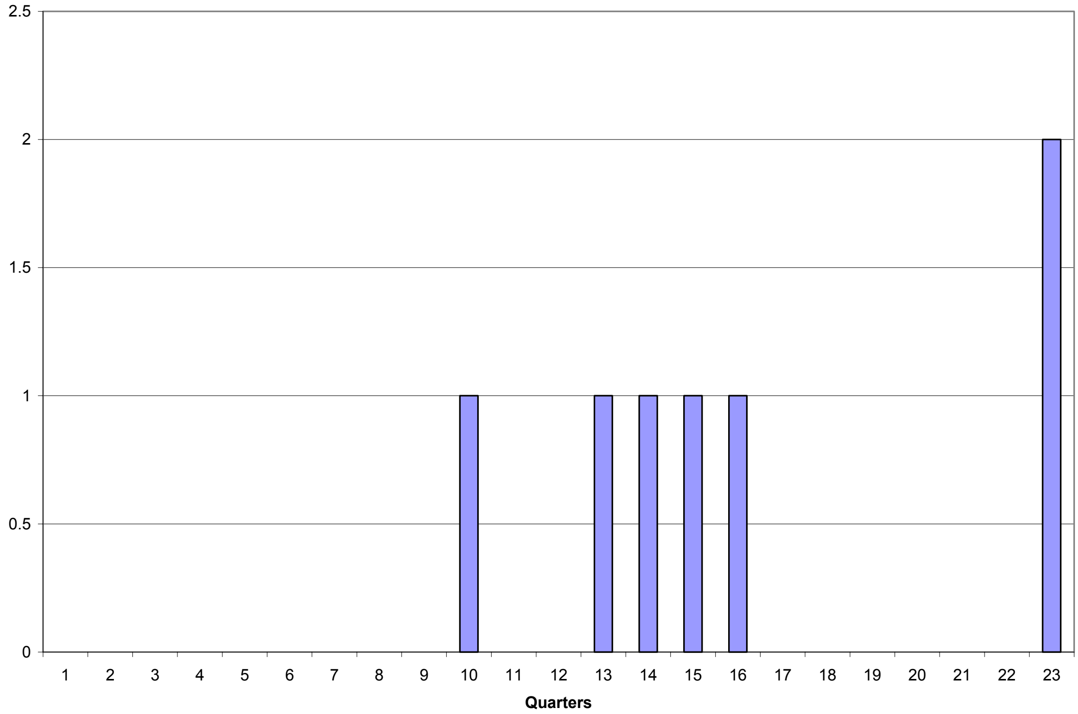

Autoregressive time series do not have a long memory – they do not remember how long it has been since the last peak or trough, leading to a statistical distribution of the length of each cycle (measured as peak to peak) that looks rather like a Poisson or Gamma distribution. Quasi-cyclic data series have cycle lengths that are clustered tighter around particular values than one would expect from a process with no memory. Australian building approvals have cycle lengths clustered around 15 and 23 (the cycle length of 10 was due to the extraordinary influence of the introduction of GST).

When the cycle lengths are very tightly clumped around a single value, this may indicate a seasonal process rather than a quasi-cyclic process. In these cases one should try to use a simple and tractable seasonal model first, rather than a more complex regime switching model.

Graph 4: Stagnant Periods - Building Approvals

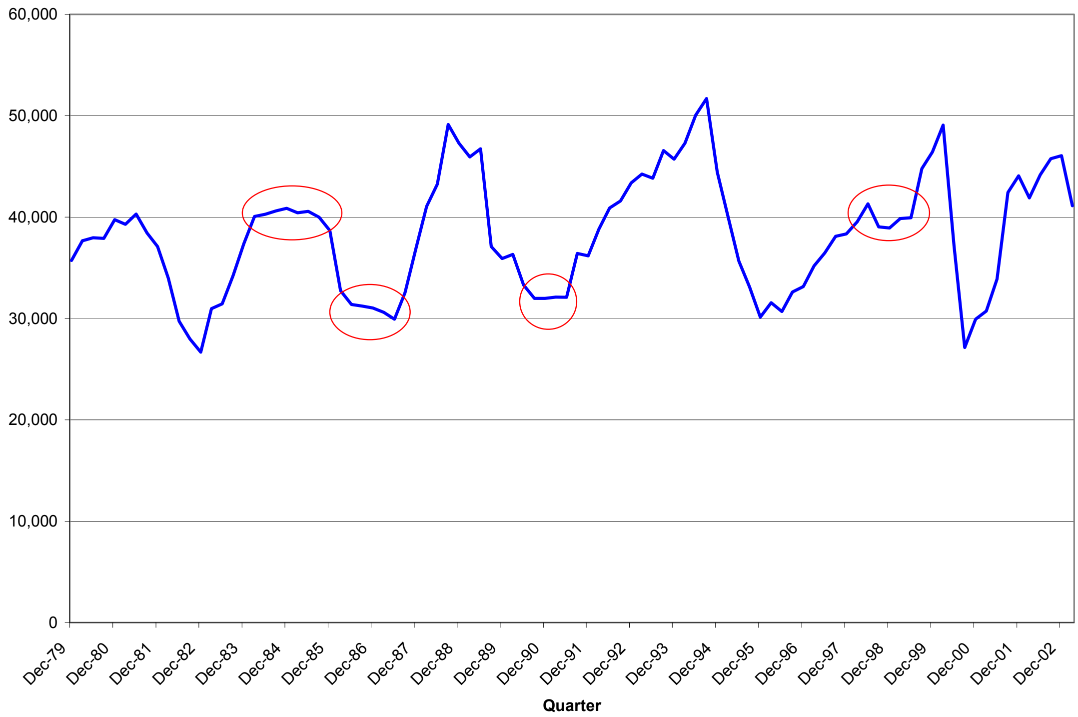

Some regime switch processes have "attractors" that periodically change. Once the model has reached the attractor value, the observed values do not show significant movement until the attractor location changes. Australian building approvals show some signs of this trait, but we will come across a clearer example of stagnant values later in this paper.

Finally, quasi-cyclic data series often test positive for a "unit root". This means that when one fits an autoregressive model to the data, the parameters may lead to a time series that is not stationary. Stationary series always tend back towards their mean values (because they gradually dampen the effect of noise), while series that are not stationary can permanently move away from the mean. The presence of a unit root can sometimes mean that the series is a random walk rather than an autoregressive series.

Some Interesting Economic Data Series

I have chosen three particular economic data series to look at because they all effect mortgage insurance claim costs, and because I had a prior expectation that they are all quasi-cyclic.

Unemployment

Unemployment can trigger a mortgage insurance claim because the loss of household income can cause the borrower to default on their home loan repayments.

Graph 5: Unemployment Rates

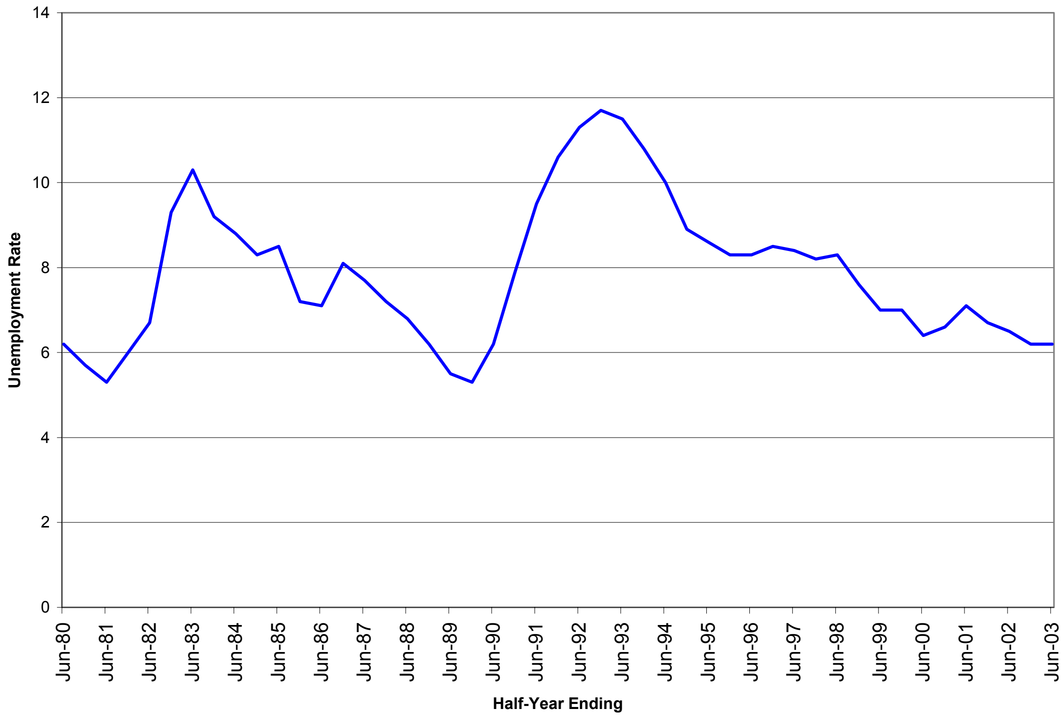

Changes in unemployment rates are strongly correlated with economic growth, with strong economic growth reducing unemployment rates. Economic growth is thought to be quasi-cyclic. Low unemployment rates can mean shortages of workers, which can in turn lead to inflationary pressures and constraints upon production, which in turn can eventually lead to slower economic growth and higher unemployment. This flow-on to economic growth could be viewed as a quasi-cyclic process.

House Prices

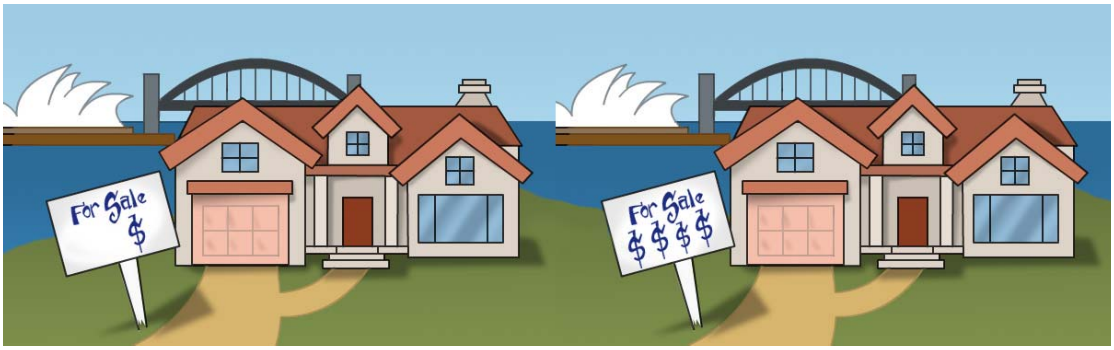

House prices affect mortgage insurance claims because the house can be sold to recover some or all of the outstanding loan balance. In times of rapid escalation in home prices, the sale of the house will usually cover the outstanding loan balance. But when home prices drop or are stagnant, the outstanding home loan balance may exceed the sale price of the house (after sales expenses e.g. real estate agent fees).

Graph 6: House Prices in My Suburb

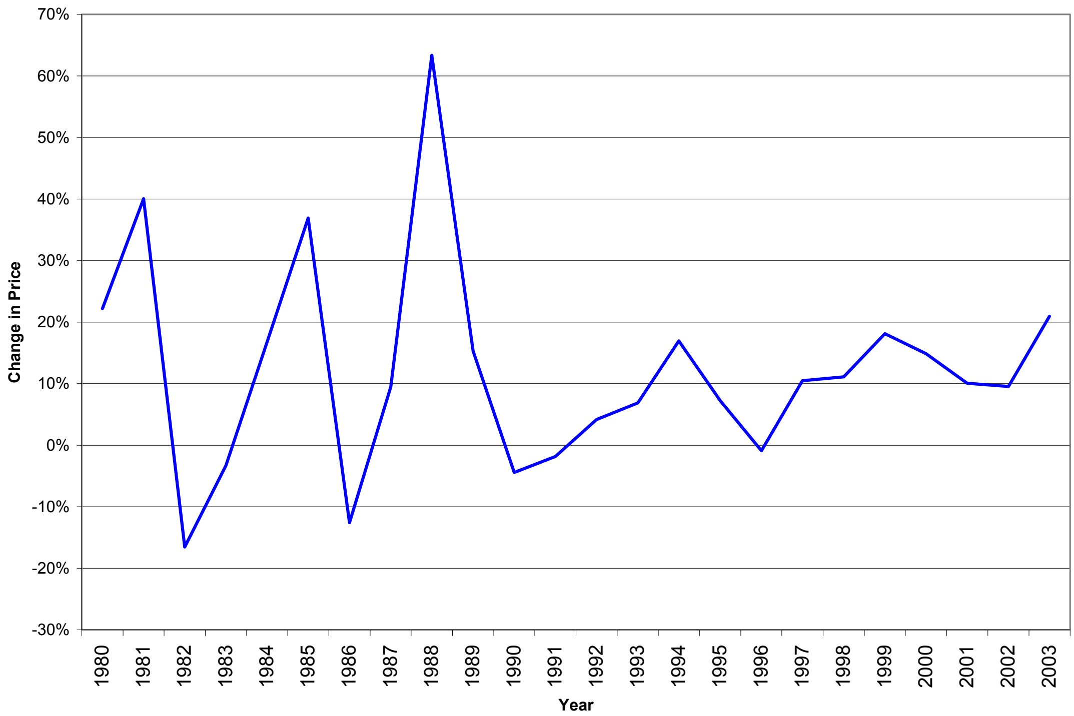

House prices can have “bubbles” (much like the dotcom sharemarket bubble), and it appears that we are in a house price bubble at the moment. Bubbles are fed by expectations and cash flows and tend to be quasi-cyclic.

Interest Rates

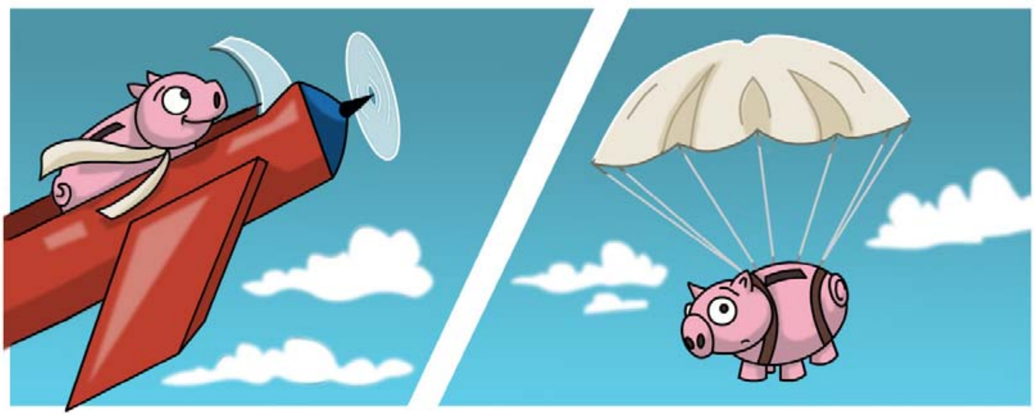

Increases in interest rates can trigger mortgage insurance claims because they can cause loan repayments to exceed what the borrower can afford to pay.

Graph 7: Overnight Cash Rates

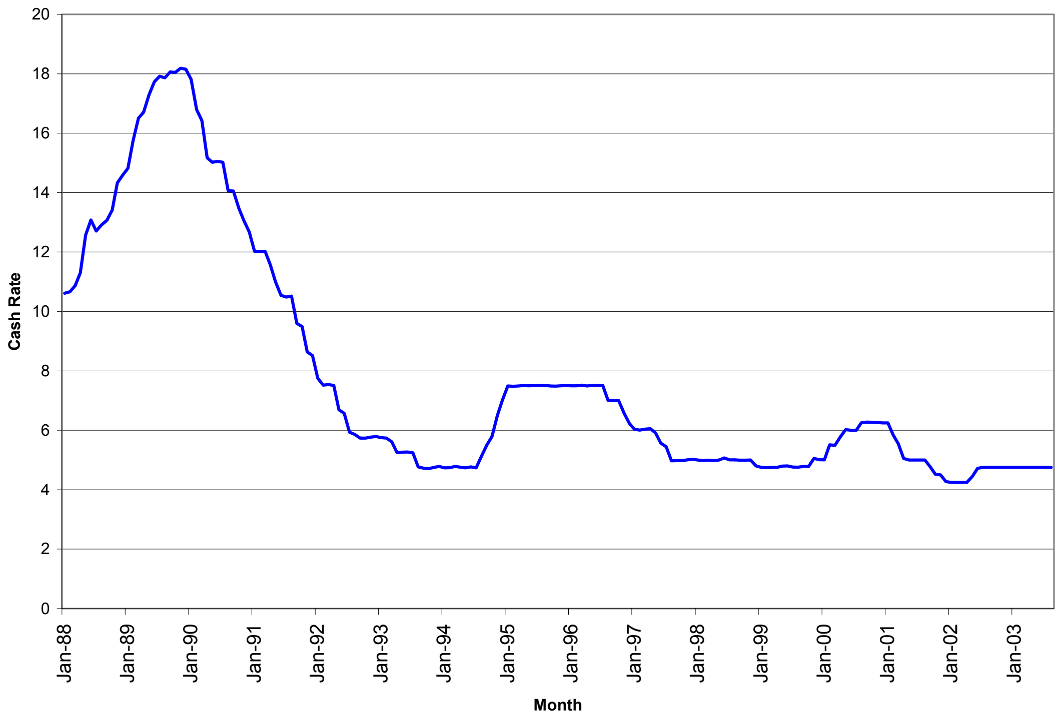

Interest rates are used by the Reserve Bank to tune economic growth and control inflation. Economic growth is thought to be quasi-cyclic.

Testing for Quasi-Cyclic Characteristics

Unit Roots

Statistical tests of the null hypothesis that a time series is non-stationary against the alternative that it is stationary are called “unit root” tests. The term “unit root” derives from the fact that an ARMA process is non-stationary if the characteristic polynomial has a root that does not lie within the unit circle of complex numbers.

An AR(1) model can be expressed as:

Equation 1: Formula for AR(1)

$$Y_t - \mu = \rho * (Y_{t-1} - \mu) + e_t$$

A unit root can occur in an AR(1) model when the reversion factor ρ equals one. In such cases the time series is not stationary and is more of a random walk than a mean reversionary series.In such circumstances the impact of a particular shock e_t does not diminish with age.

For the higher order autoregressive models, the model AR(n) can be expressed as:

Equation 2: Formula for AR(n)

$$Y_t - \mu = \alpha_1 * (Y_{t-1} - \mu) + \alpha_2 * (Y_{t-2} - \mu) + \ldots + \alpha_n * (Y_{t-n} - \mu) + e_t$$

The characteristic polynomial for this model is expressed as:

$$m^n - \alpha_1 * m^{n-1} - \alpha_2 * m^{n-2} - \ldots - \alpha_n$$

If m = 1 is a root of this polynomial, then

$$1 - \alpha_1 - \alpha_2 - \ldots - \alpha_n = 0$$

It therefore follows that:

$$\mu * (1 - \alpha_1 - \alpha_2 - \ldots - \alpha_n) = 0$$

This implies that μ is not identifiable in equation 2 above whenever there is a unit root (i.e. whenever m = 1 in the characteristic equation).

Sherris, Tedesco and Zehnwirth (1997) state that Australian interest rates may have a unit root.

There is a choice of statistical tests for unit roots, all of which have the existence of a unit root as the null hypothesis that the test then seeks to disprove this hypothesis. In this paper I have chosen to use the Dickey-Fuller (Case 2: Constant Term but No Time Trend Included in the Regression; True Process is a Random Walk) statistical test as recommended by Hamilton (1994). In this test one expresses the time series in the form:

Figure 1: Alternate AR(1) Formula

$$Y_t = \alpha + \beta * Y_{t-1} + e_t$$

Then one tests whether β is significant using a test value that is a function of the least squares estimator of β and the number of observations.

Figure 2: Dickey-Fuller Test Statistic

$$T = n * (B - 1)$$

where B is the sample estimator of β, n is the number of observations and T is the sample test statistic. The critical values for this test statistic were tabulated by Fuller (1976). The null hypothesis (unit root) is rejected when the sample statistic is less than the critical value.

Table 1: Test Results for Unit Roots

Data Series	Sample Statistic	Critical Value	Conclusion
Unemployment	-4.9	-16.8	May have a unit root
House Prices	-23.7	-16.8	May have a unit root
Interest Rates	1.1	-16.8	May have a unit root

In each of the tests I was unable to reject the hypothesis that the series has a unit root. Since the null hypothesis was that a unit root existed, these tests do not prove the existence of a unit root, and do not disprove the assumption of an autoregressive time series. However, they highlight the danger of applying an autoregressive model to this data.

Length of Runs and Autocorrelation of Deltas

In order to reject autoregressive models, one must use a statistical test that has an autoregressive process as the null hypothesis, and which is able to disprove that hypothesis.

The form of this test was suggested by a graphical presentation of the official cash rate time series. The official cash rate data appears to be partially cyclic, whereas a standard AR(1) process will not appear cyclic. The reason that the data series appear cyclic may be because the *changes* in interest rates from month to month may be correlated.

In an AR(1) autoregressive series, the values taken in consecutive time periods are correlated, but the single period changes in values (which I call "deltas") will be only insignificantly correlated at most.

We can test the sample correlation coefficient of the changes in data series values against the probability that such a sample value would occur if the time series was AR(1) with best fit parameters. I have use Monte Carlo simulation to estimate the probabilities of the sample correlation of changes in data series values being as different from the mean of the changes in data series values if they followed an AR(1) process.

Table 2: Test Results for Delta Autocorrelations Against AR(1) Process

Data Series	Sample Statistic	Critical Values	Conclusion
Unemployment	0.51	-0.34 and 0.22	AR(1) rejected
House Prices	-0.28	-0.75 and -0.18	AR(1) possible
Interest Rates	0.47	-0.14 and 0.13	AR(1) rejected

For two of the data series we are able to reject the null hypothesis that the series followed an AR(1) process. Changes in values from period to period are too strongly correlated with one another. For the case of house prices the data series may be too short to draw a conclusion. Regardless of the inability to reject the AR(1) hypothesis for house prices using delta autocorrelations, the null hypothesis was already thrown into question by the results of the unit root test in the previous section.

The AR(1) bootstrapping approach was first proposed by Priest (2002). Its weakness is the narrowly defined null hypothesis of an AR(1) process i.e. that the lag is no greater than 1 period. In order to broaden this test, we can extend the null hypothesis to an AR(n) process with a greater lag. So for the extended test we keep increasing the lag (n) until we have reached the maximum lag for which the parameters remain statistically significant.

Table 3: Test Results for Delta Autocorrelations Against AR(n) Process

Data Series	n	Periods	Sample Statistic	Critical Values	Conclusion
Unemployment	8	47	0.51	0.28 and 0.78	AR(8) possible
House Prices	11	24	-0.28	-1 and 1	AR(11) possible
Interest Rates	30	188	0.47	-0.07 and 0.44	AR(30) rejected

While we can not reject the AR(n) null hypothesis for unemployment and house prices, we need to use a large number of parameters within the longer lag autoregressive time series. The higher number of parameters may mean that we need to overparameterise the model in order to match the characteristics of the data. This should warn the user that an AR(n) model is not appropriate model to use. The house prices AR(n) model which uses 13 parameters (11 lags, a mean and a sigma) to describe 24 data points is an obvious candidate for the label of “overparameterised”.

We can define a “run” as the number of consecutive periods during which the data series values increase or decrease. This will be related to the delta autocorrelations because the higher the delta autocorrelations, the longer the runs will be. Like the delta autocorrelations, we could use a Monte Carlo method to simulate an AR(n) process to measure the distribution of lengths of runs and compare this to the observed run lengths. Because of the overparameterisation displayed above, I have used an AR(1) process for the comparison of observed and expected run lengths.

Graph 8: Unemployment Run Lengths

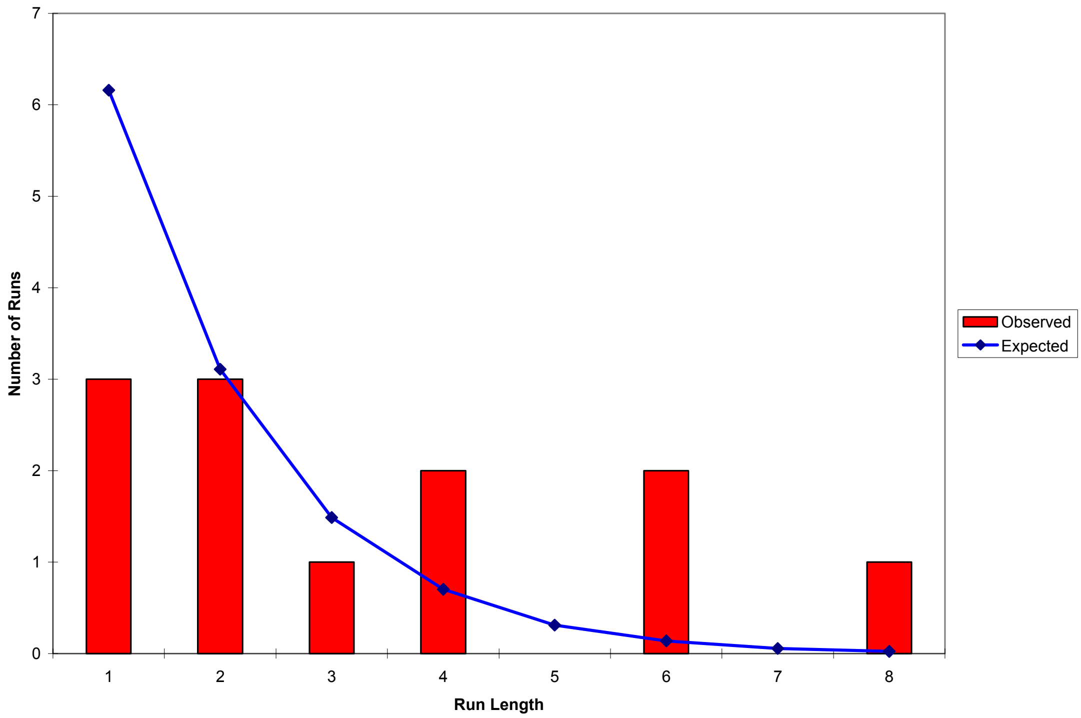

The run lengths for unemployment are much longer than would be expected from an AR(1) process. The runs of 1 and 2 periods length are during periods of stagnancy, which are discussed later in this paper.

Graph 9: House Prices Run Lengths

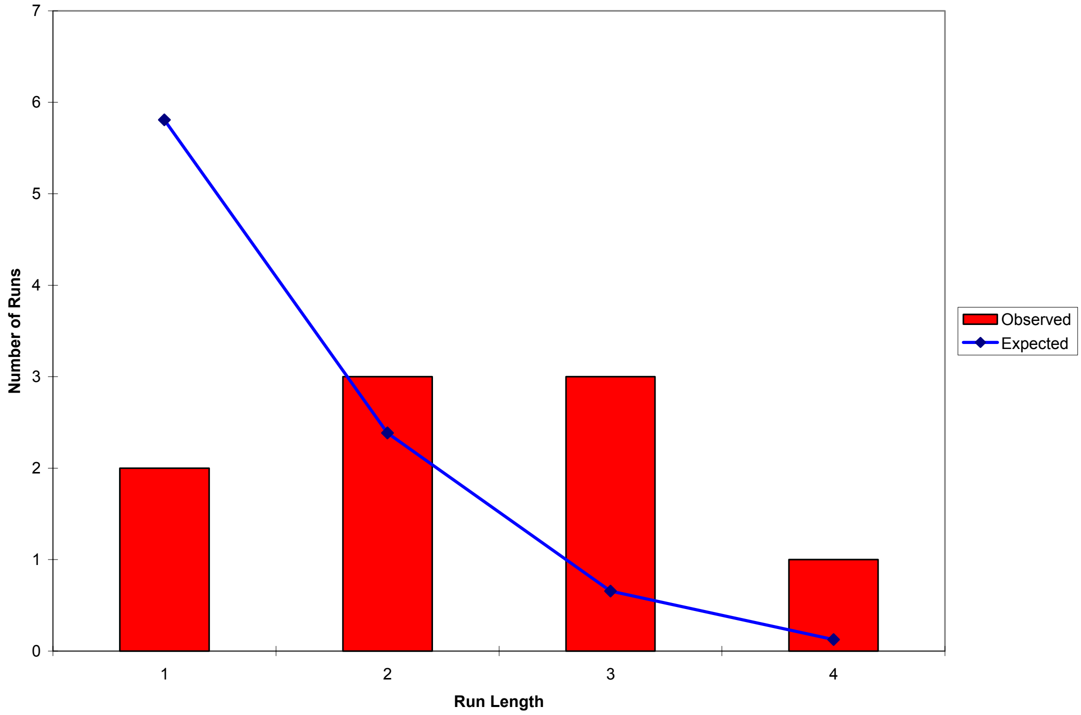

House price run lengths are more symmetrical that one would expect from an AR(1) process. This may indicate clumping of cycle lengths, which would indicate that booms in house prices occur on a quasi-regular basis.

Graph 10: Interest Rate Run Lengths

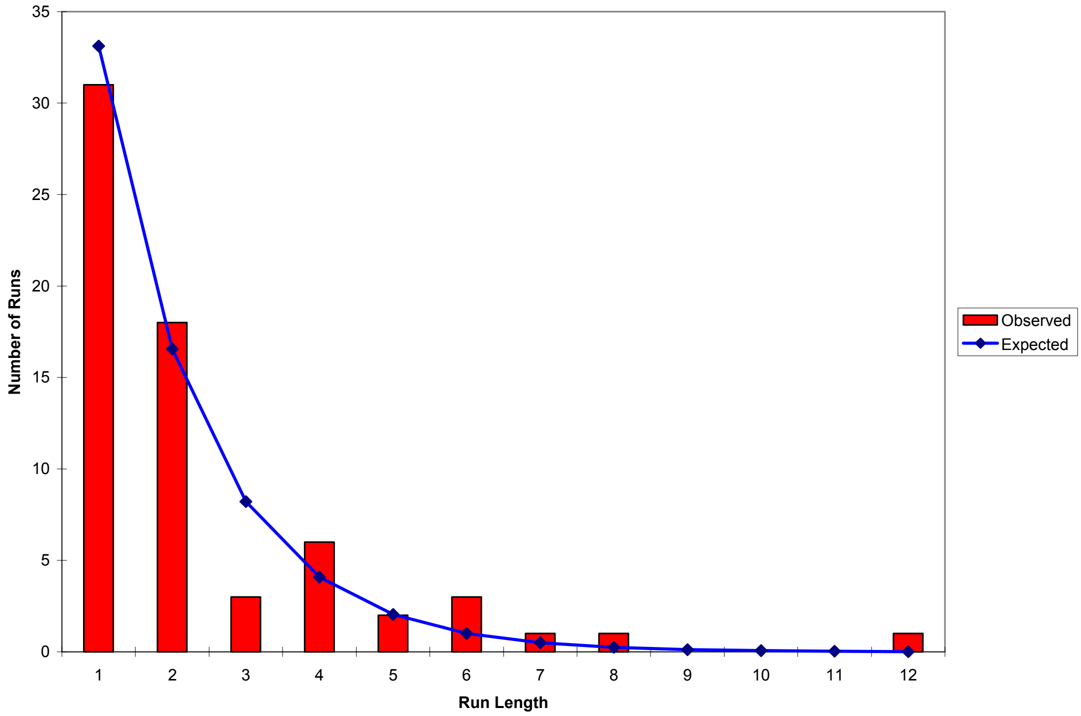

Interest rate run lengths are much closer to the statistical distribution (under the hypothesis of an AR(1) process) than we saw in the previous two examples.

Clusters of Peak / Trough Heights and Attractor Points

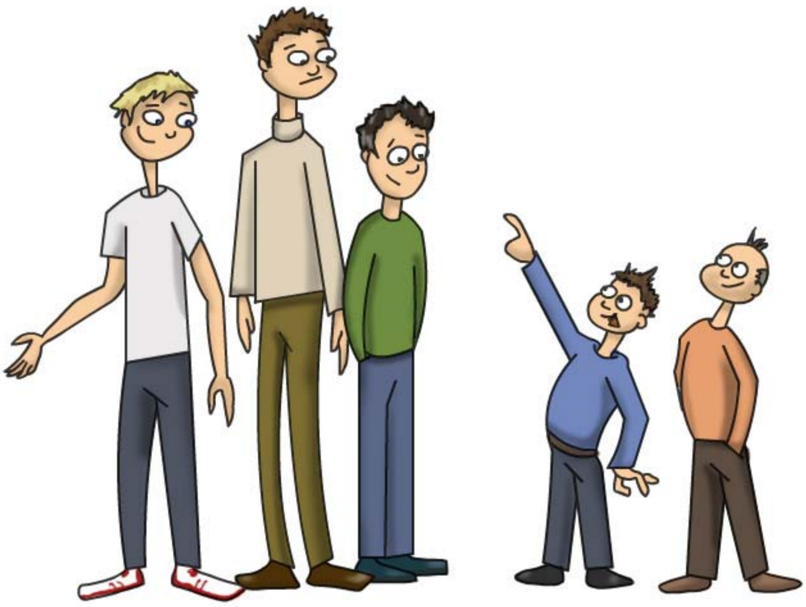

When peaks and troughs tend to recur at similar levels over time, a quasi-cyclic model with multiple "attractors" may be the most appropriate model to use. Periods of stagnancy may also occur around the attractor locations. Autoregressive time series models have only one attractor – the series mean.

We can define a peak or trough as a location at which a run ends. Alternatively we can define the peak or trough as a local minimum or maximum.

Graph 11: Attractors - Unemployment Peak and Trough Locations

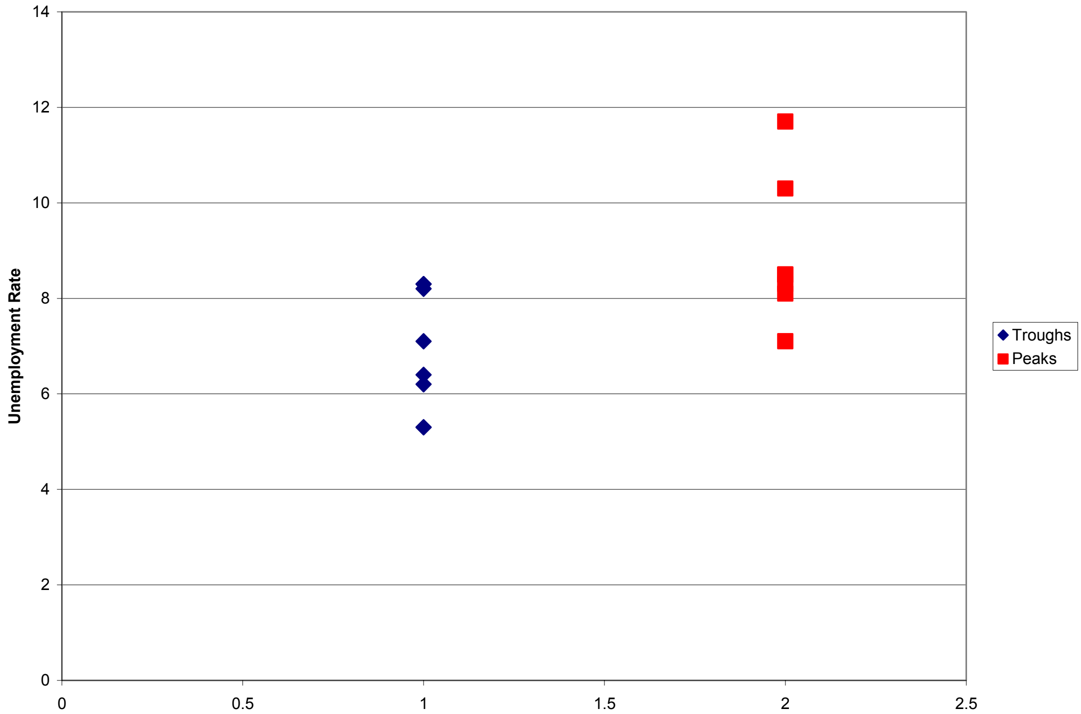

The peak and trough points for unemployment appear to have common attractor values. The strongest attractor is around 8.3 to 8.5, which contains 6 of the 15 turning points.

House price movements do not have any clear attractors.

3 of the 7 interest rate peaks and troughs lie at 4.75%. However, the sample size is too small (because the run lengths are too long) for any strong conclusions to be made.

Graph 12: Attractors - Interest Rate Stagnancy Locations

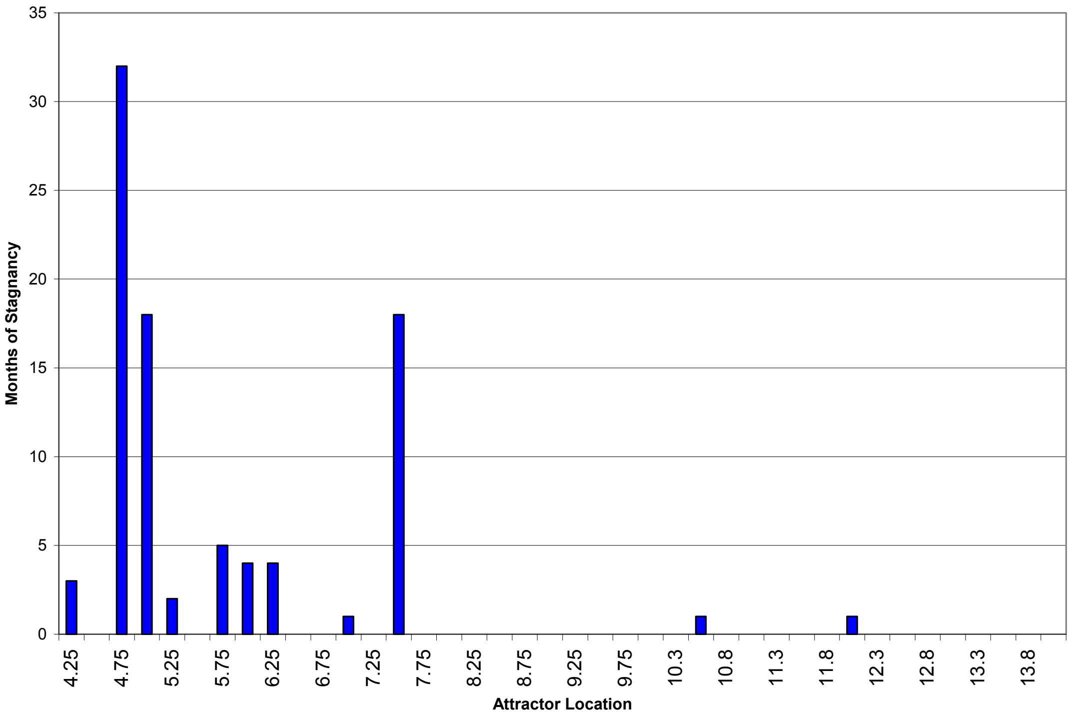

The stagnancy locations are clustered around 4.75% and 6.0%. These are also the locations of 5 of the 7 peaks and troughs in the data series. Interest rates move towards attractors that periodically change (Priest 2002).

Conclusion

While this paper presents some statistical tests that can be useful to determine whether a data series is autoregressive, some of these tests are relatively weak. The actuary should look at all of the characteristics of a data series, and consider the causes of any patterns, before using their judgement to determine whether to model using the more traditional autoregressive time series models or switching to one of the quasi-cyclic processes (such as a regime switching process).

Bibliography

Dickey, D., Fuller, W. “Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root.”, Econometrica 49 (1981) p1063

Fuller, A. “Introduction to Statistical Time Series”, Wiley, New York (1976) p371

Hamilton, J. “Time Series Analysis” Princeton University Press, New Jersey (1994)

Priest, C. “Regime Switching and Cycles”, IAA General Insurance Seminar 2001

Priest, C. “Interest Rates are Not Mean Reversionary”, International Conference on Applied Statistics, Actuarial Science and Financial Mathematics 2002

Sherris, M., Tedesco, L. and Zehnwirth, B. "Stochastic Investment Models: Unit Roots, Cointegration, State Space and GARCH Models.” Actuarial Research Clearing House no. 1, (1997): 95-144.