Economic Forecasting
Lecture 3: Characterising Cycles

Richard G. Pierse

1 Introduction

Cycles are present in many economic time series. In macroeconomics, variables such as output and employment are subject to irregular business cycles. Various economic theories have been proposed to account for these: *trade cycle theory* in the 1950s and *real business cycle theory* since the 1980s. Agricultural production, especially in developing countries, is subject to climate and weather cycles, caused by varying levels of rainfall, droughts and floods. Government policy variables may be subject to political cycles in which government policy follows a cyclical pattern determined by proximity to the next general election. Cyclical dynamics may be complicated and cyclical patterns may be more or less regular.

There are two main ways of characterising a cyclical time series, known as the *time domain* and the *frequency domain*. In the *time domain*, the properties of the series are characterised by looking at the *autocorrelation function* and the *partial autocorrelation function*, which show the correlation between observations of the series at different points in time. In the *frequency domain*, the properties of the series are characterised by looking at the *spectral density function*, which shows the composition of the series in terms of cycles of different frequencies (or equivalently periods). Both time domain and frequency domain are equally valid ways of looking at a cyclical time series. For most people, the time domain seems more natural and has the advantage of being mathematically rather simpler. For this reason, we will concentrate on time domain analysis of cycles in these lectures. However, at some points it will be useful to consider the equivalent frequency domain interpretation of the time domain concepts that will be introduced.

This lecture is a preparation for the consideration next week of the most influential forecasting model of cycles in the time domain, the *Box and Jenkins* or *ARMA* model.

2 Simple Cycles

2.1 A deterministic cycle: the cosine wave

The simplest model of a regular cycle is the cosine wave

$$C_t = a\cos(bt) \tag{1}$$

with parameters a and b. The parameter a represents the size or *amplitude* of the wave or cycle and b is the *frequency* of the cycle and is related by

$$b = \frac{2\pi}{q}$$

to q which is called the *period* of the cycle and gives the number of time periods of one complete cycle.Figure 1 represents a cosine wave with amplitude 1 and period

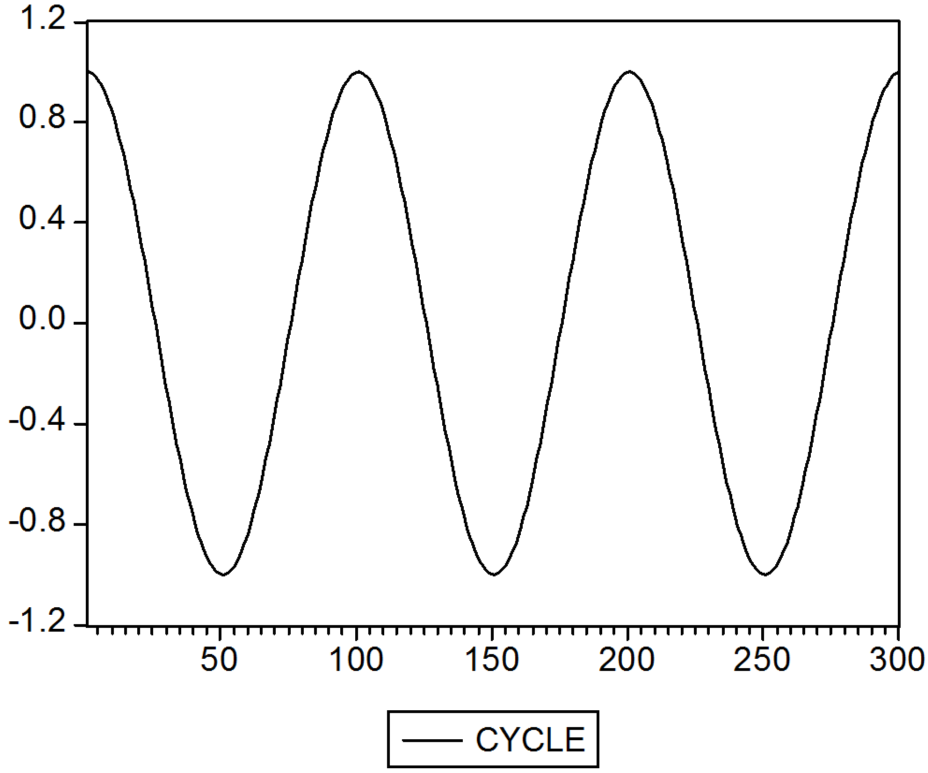

Figure 1: Cosine function with $a = 1$ and $b = 2\pi/100$

100.

In the simple cosine wave (1), the amplitude of the wave a is constant. This is what is called an *undamped cycle*. It is also possible to define a *damped cycle*

$$C_t = r^t a \cos(bt) \tag{2}$$

where r is a *damping factor* with $|r| < 1$.In this model, the amplitude of the cycle

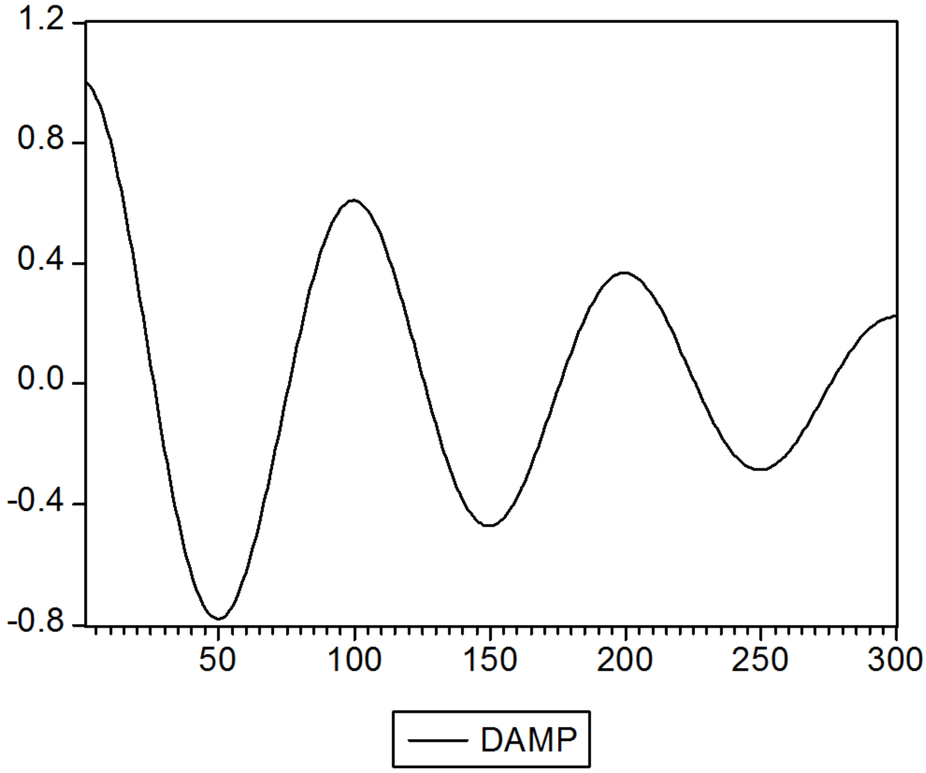

Figure 2: Damped cosine function with $r = 0.995$, $a = 1$ and $b = 2\pi/100$

decreases with time at a rate determined by r and the wave will eventually die away completely.

2.2 A stochastic cycle: the AR(2) model

An alternative stochastic model of a cycle can be represented by the equation

$$y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \varepsilon_t \tag{3}$$

where ε_t is a '*white noise*' random disturbance satisfying

$$\mathrm{E}(\varepsilon_t) = 0,\ \mathrm{E}(\varepsilon_t^2) = \sigma^2,\ \mathrm{E}(\varepsilon_t \varepsilon_s) = 0,\ s \neq t.$$

This model is known as a *second order autoregressive* model or *AR(2)* model. We will be looking at the properties of autoregressive models in more detail in the next lecture. For now we note that this model gives rise to a damped cycle when the parameters ϕ_1 and ϕ_2 satisfy

$$\phi_1 = 2r\cos(b)$$

and

$$\phi_2 = -r^2$$

with $|r| < 1$. Note that, since $-1 \leq \cos(b) \leq 1$, it follows that $-2 < \phi_1 < 2$ and $-1 < \phi_2 < 0$. We will see next week that this corresponds to the case of a pair of *complex conjugate roots*. The *amplitude* of the cycle generated by (3) is determined by the initial values y_0 and y_{-1}.

In the absence of any shocks (with $\sigma^2 = 0$), this model generates a damped cosine wave identical to the damped cosine model (2). However, when shocks are allowed to influence the process, the cycle generated becomes very irregular.Figure 3 shows the big effect of introducing even a small sized shock $\sigma^2 = 0.000025$ into the model (3). The cycle now has an irregular period and amplitude. Compare this with Figure 4 which shows the deterministic damped cosine cycle model

$$y_t = r^t a\cos(bt) + \varepsilon_t$$

where an additive disturbance has been introduced. Even with a much larger disturbance variance, $\sigma^2 = 0.01$, the cycle still exhibits a regular period and a smoothly damped amplitude.

3 Stationarity

In characterising cycles, we restrict ourselves to the consideration of *stationary time series*. This means that, if the original series is non-stationary, then it is assumed that the non-stationary trend has been removed by appropriate detrending as considered last week.

Recall from last week the definition of *weak stationarity* or *covariance stationarity*. A time series $\{y_t\}$ is said to be *covariance stationary* if it satisfies the following three conditions:

$$\mathrm{E}(y_t) = \mu \tag{4}$$

$$\mathrm{var}(y_t) \equiv \mathrm{E}(y_t - \mu)^2 = \sigma^2 < \infty \tag{5}$$

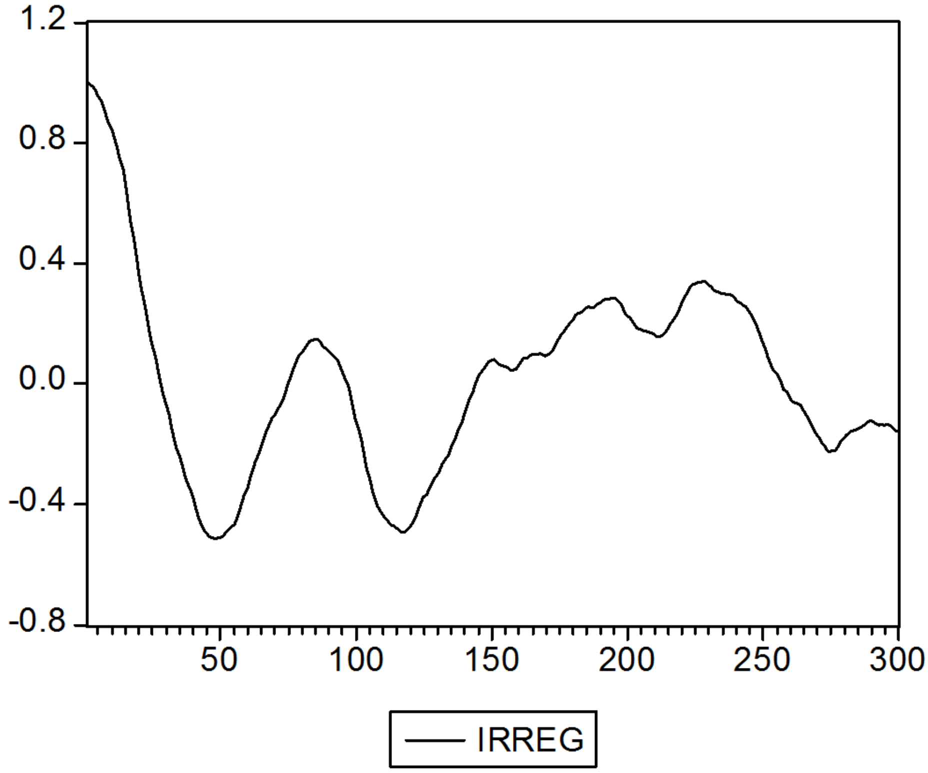

Figure 3: AR(2) model with $r = 0.995$, $b = 2\pi/100$, $\sigma^2 = 0.000025$

and

$$\text{cov}(y_t, y_{t-s}) \equiv \text{E}(y_t - \mu)(y_{t-s} - \mu) = \gamma_s. \tag{6}$$

The first condition (4) states that the mean of the series is the same in every time period t. The second condition states that the *variance* of the series is also the same in every time period t. The third condition states that the *covariance* between the two observations y_t and y_{t-s} depends only on the *distance* between them, s, and is the same in every time period t.

Note that the variance is a special case of the covariance where $s = 0$ so that we may write

$$\text{var}(y_t) = \gamma_0.$$

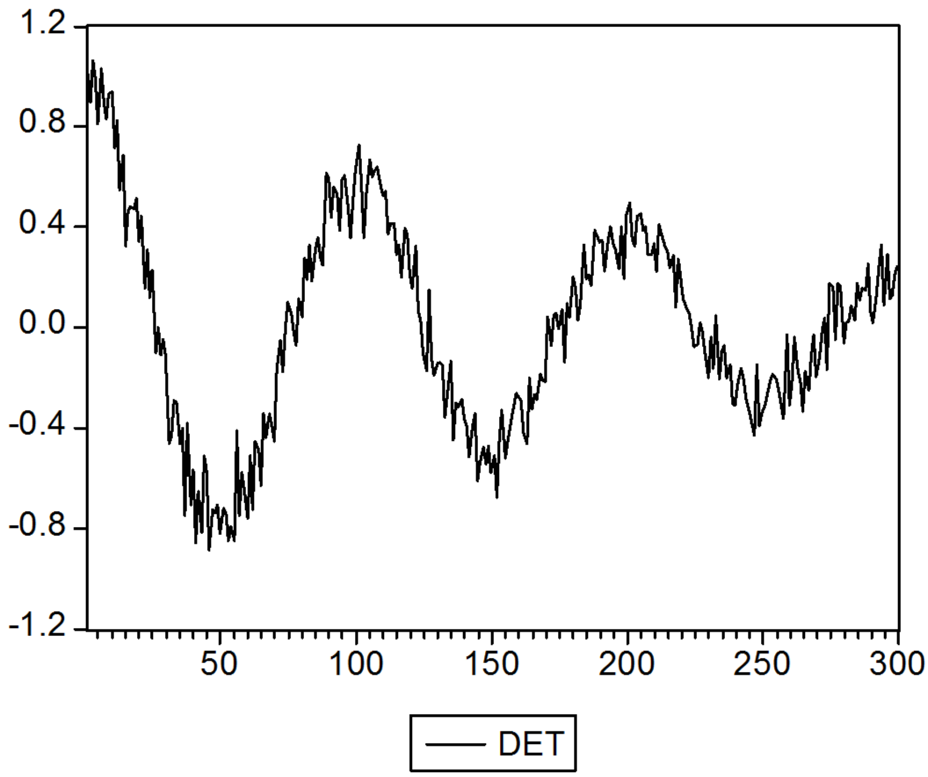

Figure 4: Damped cosine function with disturbance, $\sigma^2 = 0.01$

4 The Autocorrelation Function (*ACF*)

For a *covariance stationary* series, the covariance depends only on the distance between observations or *displacement*. This means that we can consider the covariances as a function of s, known as the *autocovariance function*

$$\gamma(s).$$

Note that this function is symmetric with $\gamma(s) = \gamma(-s)$ since

$$\gamma_{-s} = \text{cov}(y_t, y_{t+s}) = \text{cov}(y_{t-s}, y_t) = \gamma_s.$$

The autocovariance function has a disadvantage in that it is not scale independent but depends on the units in which y_t and y_{t-s} are measured. Instead, we work with the *autocorrelation function* or *ACF*

$$\rho(s) = \frac{\gamma(s)}{\gamma(0)} = \frac{\text{cov}(y_t, y_{t-s})}{\text{var}(y_t)}$$

which gives the simple correlation between observations y_t and y_{t-s}. Note that the autocorrelation function is also symmetric with $\rho(s) = \rho(-s)$ and has the additional property that $\rho(0) = 1$ since an observation is always perfectly correlated with itself. Also, a correlation must lie between minus 1 and plus one so that

$$-1 \leq \rho_s \leq 1, \quad s = 0, \pm 1, \pm 2, \cdots$$

For a stationary series it must be the case that eventually the autocorrelations disappear so that, as $s \to \infty$, $\rho_s \to 0$.

The autocorrelation function gives a complete characterisation of a stationary time series. The shape of the *ACF* shows how the autocorrelations behave as the distance between observations increases. Eventually, we know that, with any stationary series, the autocorrelations have to go to zero. However, whether they cut off abruptly or die away gradually or oscillate will characterise the series.

5 Estimating the *ACF*: the Correlogram

The autocorrelation function $\rho(s)$ is defined by the theoretical expression

$$\rho(s) = \frac{\mathrm{E}(y_t - \mu)(y_{t-s} - \mu)}{\mathrm{E}(y_t - \mu)^2}.$$

In order to make this operational, we need to be able to estimate the function from a sample of T observations for y_t, $t = 1, \cdots, T$.

The *sample autocorrelation function* or *correlogram* is a *consistent* estimator of the *ACF* and is defined by

$$\widehat{\rho}(s) = \frac{\frac{1}{T}\sum_{t=s+1}^{T}(y_t - \overline{y})(y_{t-s} - \overline{y})}{\frac{1}{T}\sum_{t=1}^{T}(y_t - \overline{y})^2} \tag{7}$$

where

$$\overline{y} = \frac{1}{T}\sum_{t=1}^{T} y_t$$

is an estimate of the sample mean μ. Note that the summation in the numerator is over $T - s$ observations but is divided by T. Asymptotically, this makes no difference to the properties of the estimator.

6 The Box-Pierce and Box-Ljung Q Statistics

It is useful to be able to test the hypothesis that all the autocorrelations up to a certain order m are equal to zero. This is a test of the hypothesis that the series

y_t is 'white noise' with $\text{cov}(y_t, y_{t-s}) = 0$, $s > 0$. Box and Pierce (1970) proposed a statistic based on the sum of squared estimated autocorrelations:

$$Q_{BP} = T \sum_{s=1}^{m} \widehat{\rho}_s^2 \sim \chi_m^2.$$

This statistic asymptotically has a chi-squared distribution with m degrees of freedom on the null hypthesis that all the autocorrelations are zero. This test is sometimes known as a *portmanteau* test. Note that if any one autocorrelation is non-zero, then the null hypothesis is violated and we expect that the test will reject the null. However, in practice a single non-zero autocorrelation may be swamped by many zero autocorrelations and the test may fail to reject in this case. This is an example of what is called *low power* of a test.

A modified version of the test was proposed by Ljung and Box (1978). This modified form is designed to follow more closely the chi-squared distribution in small samples. The modified statistic is defined by

$$Q_{BL} = T(T+2) \sum_{s=1}^{m} \frac{1}{T-s} \widehat{\rho}_s^2 \sim \chi_m^2.$$

In this form, the sample autocorrelations are weighted and the scaling factor adjusted. Note that, when T is large in relation to m, the two forms will be very close to each other.

7 The Partial Autocorrelation Function ($\boldsymbol{PACF}$)

The autocorrelation function measures the simple correlation between y_t and y_{t-s}. The *partial autocorrelation function* or *PACF* measures the correlation between y_t and y_{t-s} *after controlling* for the effects of y_{t-1}, y_{t-2}, $\cdots$, y_{t-s+1}. For example, in the model

$$y_t = \phi y_{t-1} + \varepsilon_t$$

y_t is not directly affected by y_{t-2} but only by y_{t-1}. However, it can be shown that the simple correlation between y_t and y_{t-2} is

$$\rho_2 = \frac{\text{cov}(y_t, y_{t-2})}{\text{var}(y_t)} = \phi^2$$

where the correlation arises because of the indirect effect of y_{t-2} on y_{t-1} and hence on y_t. By contrast the *partial autocorrelation* between y_t and y_{t-2} is zero since, after controlling for the effect of y_{t-1}, there is no additional effect from y_{t-2}.

The *PACF* is denoted

$$p(s)$$

and p_s can be estimated by the estimator $\widehat{p}_s$ in the regression defined by

$$y_t = c + p_1 y_{t-1} + p_2 y_{t-2} + p_3 y_{t-3} + \cdots + p_s y_{t-s} + \varepsilon_t.$$

The estimated *PACF* is known as the sample partial autocorrelation function or *partial correlogram.*

The *PACF* provides an alternative way to the *ACF* of viewing the autocorrelations. By examining the shape of both functions, many stochastic processes can be distinguished as will be discussed in the next lecture.

8 Examples

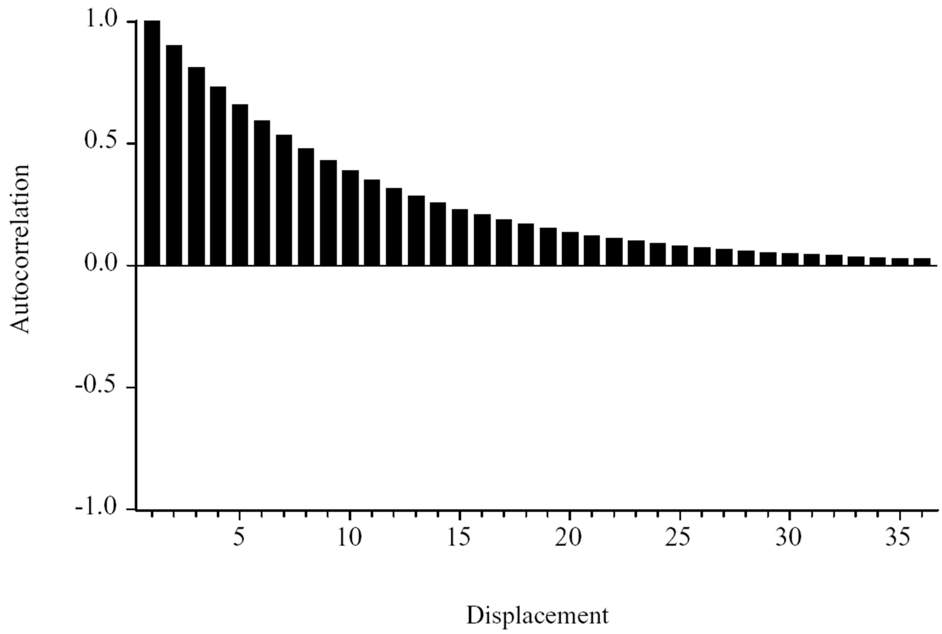

Figure 5: *ACF* with gradual damping

Figures 5–7 show autocorrelation functions with different shapes. In Figure 5 the autocorrelations die away gradually. In Figure 6 the correlations cut off sharply after lag 14 and are zero thereafter. In Figure 7 the autocorrelations follow an oscillation with gradual damping. Similar patterns may be observed in partial autocorrelation functions.

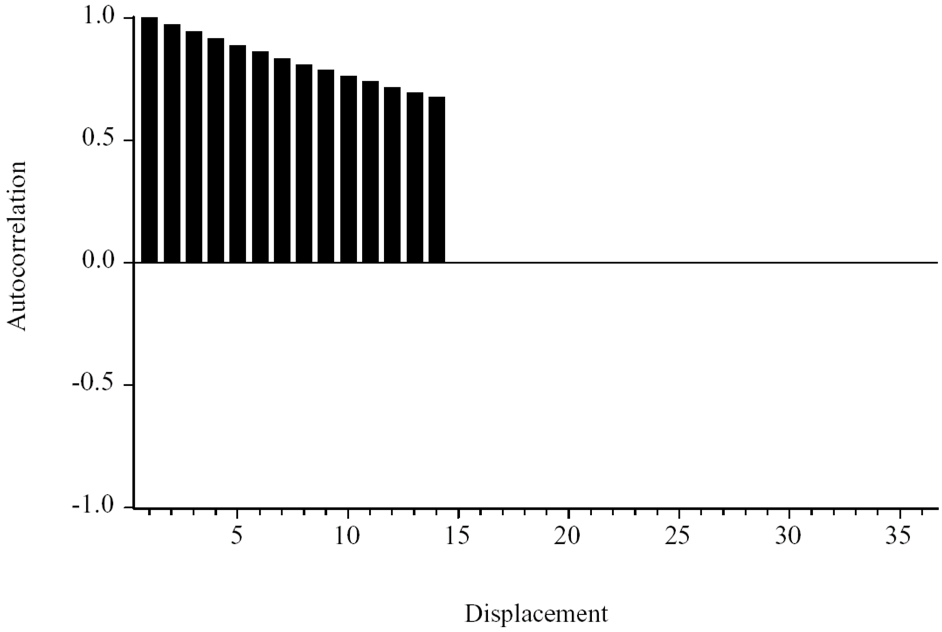

Figure 6: *ACF* with sharp cutoff

Finally, Figure 8 shows the correlogram for the *AR(2)* model of equation (3). This resembles Figure 7 in that the autocorrelations oscillate in sign but decrease in magnitude. The fact that it takes a very large number of periods for the autocorrelations to die down indicates that the series is close to being non-stationary. In contrast, the partial correlogram for this process cuts off sharply after lag 2. We will see next week that these are the characteristics of a second order *AR* process with a pair of complex conjugate roots.

References

[1] Box, G.E.P. and D.A. Pierce, (1970), 'Distribution of residual autocorrelations in autoregressive integrated moving average time series models', *Journal of the American Statistical Association*, 65, 1509–1526.

[2] Ljung, G.M. and G.E.P. Box (1978), 'On a measure of lack of fit in time series models', *Biometrika*, 65, 297–303.

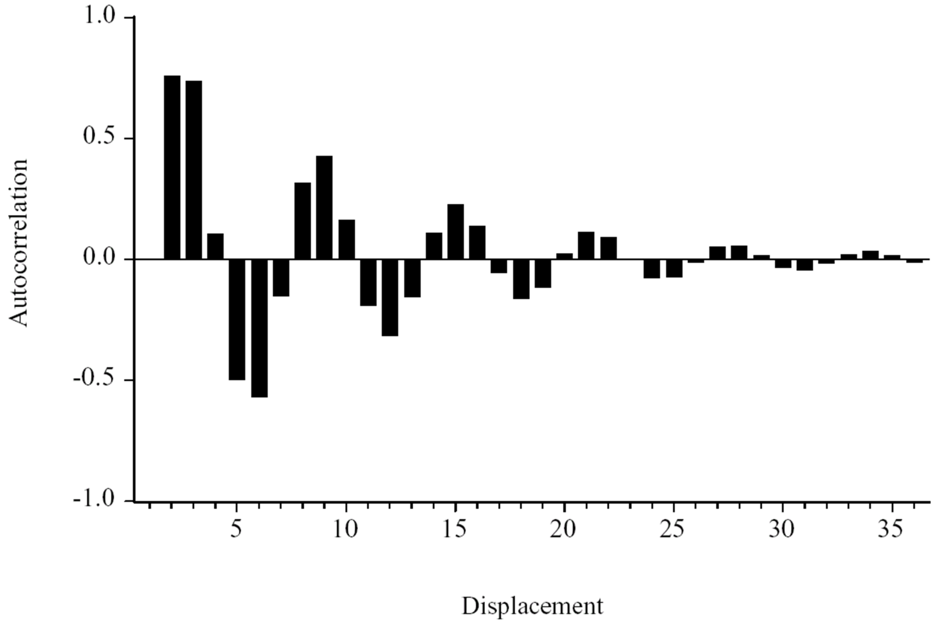

Figure 7: ACF with gradual damped oscillation

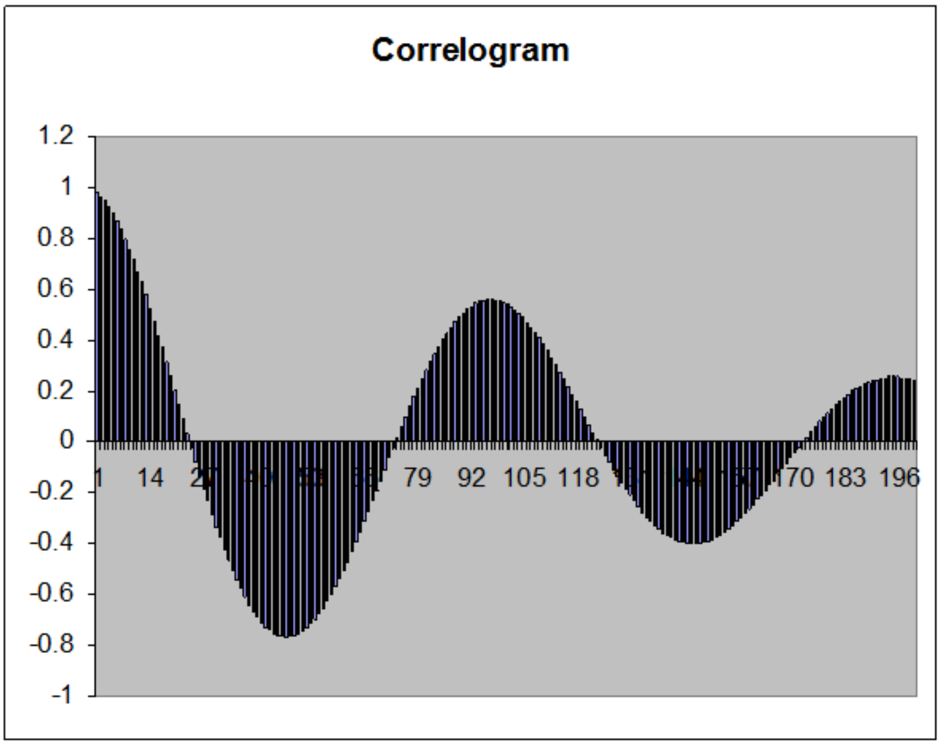

Figure 8: Correlogram for *AR(2)* model of equation (3