Forecasting Australian Unemployment Rates Using Spectral Analysis

Patrick J. Wilson and L. J. Perry [*]

University of Technology, Sydney

Abstract

Univariate spectral analysis is used to model seasonally unadjusted quarterly unemployment rate data for Australia, 1978(2) to 2002(3). Data are tested for three categories: persons, males and females. Dynamic out-of-sample forecasts are made for 8 quarters using spectral analysis models evaluated against ARIMA model counterparts. It is found that the spectral analysis models achieve higher levels of forecasting accuracy than ARIMA counterparts, including turning point forecast accuracy. These results emerge in spite of weaker in-sample explanatory power of the spectral models against the ARIMA models. It is concluded the results suggest that the spectral model is ultimately better attuned to the various cyclical forces of the past unfolding into the future.

1. Introduction

There has long been interest in modelling and forecasting unemployment rates. More recently, interest has been directed towards the modelling of apparent asymmetric properties of official unemployment rate series. The asymmetric properties of interest are the observed tendency for the unemployment rate to rise relatively rapidly and fall slowly. Univariate (i.e., single variable) analyses have been developed by Neftci (1984), Acemoglu and Scott (1994), Parker and Rothman (1997), Verbrugge (1997), Rothman (1998), Stevenson and Peat (2000), Bodman (1998, 2001), van Dijk *et al.* (2002) and Skalin and Terasvirta (2002), among many others. These studies considered data for various countries, though mainly the USA, and used a variety of techniques to capture, in one form or another, apparent nonlinearities in official unemployment rate data. Multi-variate analyses have also been developed to address the issue of non-linearities. Thus, McHugh *et al*. (2002) have developed a multi-variate model of changes in the unemployment rate incorporating the same sort of mechanism to capture nonlinearities used in a number of univariate models.

Nonlinearities in the labour market are believed to be related to a number of possible mechanisms that may or may not be present in different countries (see Siebert, 1997; Borland, 1997; Debelle and Vickery, 1998; Katz , 1998; Jackman, 1998; Chapman and Kenyon, 2002). However, labour market

rigidities in one form or another are seen to be the central factor contributing to temporal non-linearity in the unemployment rate. These labour market rigidities include an interdependent mix of wage rigidities (Valentine, 1993; Ball, 1997; Dawkins, 2002), the institutionalised prevention of the unemployed competing for the jobs of the employed (see, Blanchard, 1991; OECD, 1994) and de-motivating social welfare systems (see, Layard, *et al.*, 1991; Lindbeck,1995; Wilson, 1996; Forslund and Krueger, 1997; Burtless, 2002).

It is important to note that these sorts of rigidities can and do change. Sometimes the change occurs relatively rapidly, such as when there is a change in government with a new set of policy initiatives. Examples include new (or old) governments declaring war, or a new government implementing major changes in labour market policies, such as those pursued in Britain during the period of Prime Minister Margaret Thatcher's governance. Sometimes the changes are slower to materialise. For example, the power of trade unions in the wage setting process has declined for most OECD countries as union density has declined over the last two decades or so. At the same time, there has been a marked long-term (20 year) general decline in the reported number of industrial disputes per worker over the same time frame (see, Perry and Wilson, 2001). Added to this already complicated picture of changing short- and long-term labour market characteristics are other possible outside shocks to the labour market, such as the largely exogenous shocks of the oil-crises of the 1970s as well as the often overlooked long-term impact of declining real energy costs during much of the 1980s and 1990s.

Thus there exists over time a complicated mix of short-, medium- and long-term influences working their way through the labour market. All of these influences can impact on the official and unofficial (see, Chapman and Kenyon, 2002; Gregory, 2002) unemployment rate series. Modellers of, in particular, univariate systems seek to find mathematical regularities in the past that might be repeated into the future and thus provide an objective (and repeatable) basis for making forecasts. Arguably, those studies that test their modelling out of sample are of particular interest, as the ultimate test of any model is to see whether the regularities mathematically distilled from the past have continued into the future. In this regard the recent univariate models of Rothman (1998), Stevenson and Peat (2000) and, at least to some extent, Skalin and Terasvirta (2002) are of particular interest. All of these studies make out of sample forecasts, so let us briefly review their work.

Rothman (1998) carries out a forecasting competition comparing the out-of-sample forecasting performance of 6 non-linear univariate models. His study analyses the quarterly log-linear detrended unemployment rate for the USA. His data sample is 1948:1 to 1993:1 from which he allocates thirteen years (1980:1 to 1993:1) for out-of-sample forecasting. He evaluates 4-quarter ahead and 8-quarter ahead horizons and finds in general that forecasting is improved by using non-linear rather than linear models. Stevenson and Peat (2000) analyse Australian monthly seasonally unadjusted unemployment rate data for the period 1978:2 to 1998:7. They compare the

forecast results of a non-linear model with a standard linear model. Their
non-linear model is a logistic smooth transition autoregressive (LSTAR)
model that effectively allows for rapid regime change when the
unemployment rate rises rapidly. Their linear model is a standard auto-
regressive model. Stevenson and Peat allocate 24 months for out-of-sample
forecasting purposes. This is a much shorter period than that allocated by
Rothman, but arguably makes sense when using a shorter period for
developing a model in the first place and focussing on short-range monthly
forecasts. Like Rothman they find their non-linear model does a generally
superior job at dynamic forecasting, particularly towards the later months
of their forecast period, than the linear rival. Finally, Skalin and Terasvirta
(2002) model quarterly seasonally unadjusted unemployment rate data for
a number of OECD countries, from (variably) the 1960s and 1970s to the
mid 1990s. They also use a univariate LSTAR model. They find their model
works well for a number of OECD countries, including Australia. They do
not use their model for out-of-sample forecasting, however. Rather for
illustrative purposes they extrapolate the model for Germany to around
the year 2040, which, though at this stage untestable, gives an interesting
depiction of the (regular) shape of things possibly to come.[1]

We acknowledge the apparent presence of nonlinearity in unemployment
rate data for many countries over different periods. However, we argue in
this paper that spectral analysis may also prove to be an effective vehicle
for modelling apparent nonlinearities in the unemployment rate. How?
Spectral analysis decomposes time series data into various cycles of varying
duration that, when recombined, may generate periods of apparent rapid
asymmetric growth followed by periods of relatively slow decline. These
outcomes depend upon the way in which different relatively short and long
cycles combine. The question we raise and seek to answer in this paper is:
how does spectral analysis compare with the standard ARIMA model as an
unemployment-rate forecasting vehicle? Is spectral analysis a contender
for further and future forecasting challenges? To answer the question raised,
in section 2 we briefly review the basics of the methodology of spectral
analysis employed in this paper. Then in section 3 we compare the empirical
results for the ARIMA model with the spectral analysis model. Some
concluding thoughts are offered in section 4.

2. A Review of Spectral Analysis and ARIMA Techniques

It is reasonable to categorise forecasting models as either univariate or
multivariate or, perhaps more conventionally, as time series versus
econometric (causal based) forecasting. Unfortunately this dichotomy is
not always clear cut, for example Vector Autoregression (VARs) time series
models relate times series values of one variable to those of another or others.
It has been suggested that time series models are considered superior for
short-term forecasting, while econometric models are considered superior
over the longer term (McKnees, 1982). An underlying assumption in time
series forecasting is that the set of causal factors (macroeconomic

---

[1] See also van Dijk *et al.* (2002) who also make some long-term forecasts to 2050.

fundamentals) that operated on the dependent variable in the past will exhibit similar influence in some repetitive fashion in the future. The basic idea in developing univariate forecasting models is to somehow extract a mathematical model that will pattern this behaviour.

Spectral Models and Moving Windows [2]

Nerlove (1964) and Granger (1969) were the two foremost researchers on the application of spectral techniques to economic time series. Their pioneering work has been followed by later researchers using these same techniques either in a preliminary search for information on the nature of long and short cycles, or to simply confirm or deny cyclical patterns within data series. While spectral analysis is ideally directed at data containing cycles of a fixed length, it has found a great deal of usefulness in identifying the approximate length of cycles in data which have non-periodic cycles - it has even found a fruitful use in the identification of cyclical activities amongst terrorists (Iksoon, Cauley and Sandler, 1987). More recently Levy (2002) has used spectral techniques to examine the cointegration properties between series.

The use of spectral analysis requires a change of focus from an amplitude-time domain to an amplitude-frequency domain. Thus spectral analysis commences with the assumption that any series, $\{X_t\}$, can be transformed into a set of sine and cosine waves such as:

$$X_t = \eta + A \sum_{j=1}^{n} [\cos(2\pi \frac{ft}{n}) + \sin(2\pi \frac{ft}{n})] \tag{1}$$

where η is the mean of the series, A is the amplitude, f is the frequency over a span of n observations, t is a time index ranging from 1 to N where N is the number of periods for which we have observations, the fraction (ft/n) for different values of t converts the discrete time scale of time series into a proportion of 2π and j ranges from 1 to n where n= N/2. The highest observable frequency (the Nyquist or folding frequency[3]) in the series is n/N (i.e., 0.5 cycles per time interval). High frequency dynamics (large f) are akin to short cycle processes while low frequency dynamics (small f) may be likened to long cycle processes. If we let $\frac{ft}{n} = h$ then equation (1) can be re-written more compactly as:

$$X_t = \eta + A \sum_{j=1}^{n} [\cos(\omega_j) + \sin(\omega_j)] \quad with \quad \omega_j = 2\pi h \tag{2}$$

Spectral analysis can be used to both identify and quantify apparently non-periodic short and long cycle processes. A given series $\{X_t\}$ may contain many cycles of different frequencies and amplitudes and such combinations of frequencies and amplitudes may yield cyclical patterns which appear

[2] The section below draws heavily from Wilson *et al.* (2000).
[3] Since higher frequencies are 'folded down' into the interval $[0, \pi/\Delta]$ where Δ is the sampling interval (cf. Bloomfield, 1976). In relation to equations (1) and (2) note that there are two 'sources' of aliasing. First, positive and negative frequencies ω and $-\omega$ are indistinguishable. Second, the equal spacing of time in our analysis introduces a further aliasing (cf. Hamilton, ch.6). Finally we note that every frequency not in the $[0, \pi/\Delta]$ interval has an alias in that range, termed its principal alias.

non-periodic with irregular amplitude. In fact, in such a time series it is
clear from equation (2) that each observation can be broken down into
component parts of different length cycles which, when added together
(along with an error term), comprise the observation. Taking equation (1),
for example, and letting η = 0, A=2, f=3 & 13 and N=100 we can gain a
visual impression of long and short cycles as shown in figures 1a and 1b.
Figure 1c demonstrates, for a hypothetical data set, how the re-combining
of different frequencies can yield a recomposed outcome that has an
irregular amplitude and period, as shown by the heavy unbroken line.

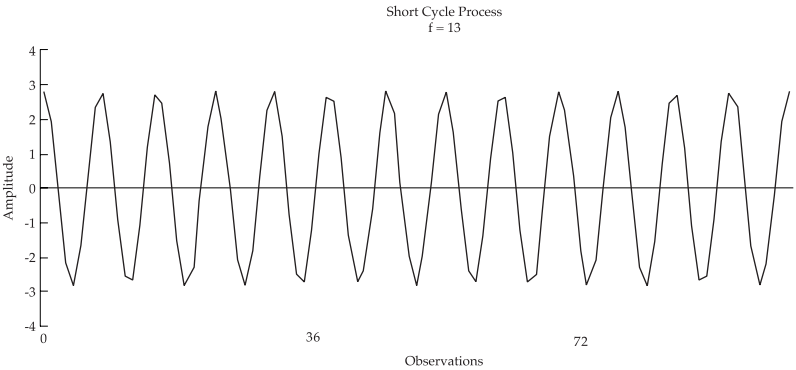**Figure 1a Thirteen Complete Cycles Over the N=100 Observations**

Figure 1b Three Cycles Over the N= 100 Observations
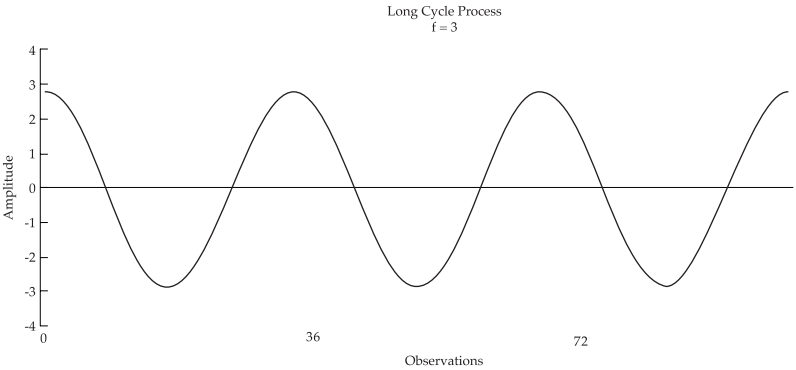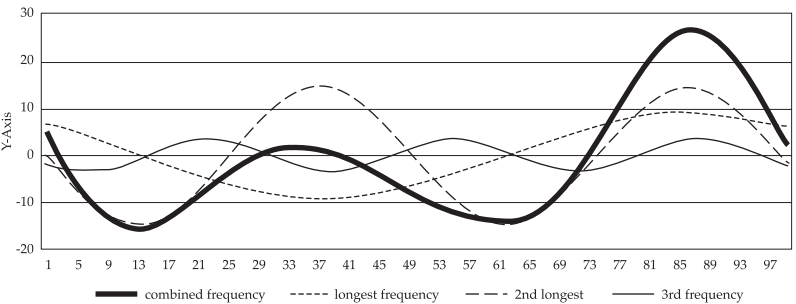

Figure 1c First Three Cyclical Components and Combined Frequencies

Spectral analysis may therefore be used to decompose and 'extract' approximate cycles of different length from the data. For purposes of forecasting, the moving spectral window approach (described below) seeks to establish a relationship between the combination of all identified cycles in one period (window) with those of another. If a method is available to project the individual cycles at different frequencies, then a re-combination of these cycles will yield a forecast (projection) of the underlying series. Furthermore, if the individual projections are such that the cycles may be in alignment, then the combined effect may be such as to indicate a (future) turning point.

Time Averaging and the Moving Window

Hannan (1963) first proposed regression analysis in the frequency domain, later examining the use of this technique in estimating distributed lag models (Hannan, 1965, 1967). Engle (1974) demonstrated that regression in the frequency domain has certain advantages over regression in the time domain. For instance he explained that the application of spectral analysis to data containing both seasonal (high frequency) and non-seasonal (low frequency) components may produce efficiencies since these different frequencies can be modelled separately and then may be re-combined to produce fitted values. The ability to decompose data series into different frequencies for separate analysis and later recomposition is the first fundamental concept in the use of spectral techniques in forecasting.

The second fundamental notion in the use of spectral analysis in forecasting is application of the time averaging algorithm developed by Welch (1967). The further advancement of this time averaging method in forecasting economics series is attributed to Ridley and Mobolurin (1987) and later Ridley (1994). Ridley and Mobolurin (1987) have called their extension the moving window method of forecasting, which can be summarised as follows: Assume the sequence $\{X_t\}$ is the depiction of a stationary stochastic process occurring at time t. Now, define a moving window of some length L in the time domain. If this window is moved forward one period at a time it will generate a new sub-set of observations in the time domain. The observation sub-set in each window can be transformed into the frequency domain as indicated earlier; therefore for each adjacent pair of windows as the length L is rolled forward there will be an observation on the input and output process for each frequency in the frequency domain. An important underlying assumption in Ridley and Mobolurin's moving window method is that the given series has a dominant frequency, as well as other frequencies that are multiples of this dominant frequency. This dominant frequency determines the length of the moving window. For instance, a business cycle may have a dominant frequency of about four years (peacetime) or six and a half years (wartime) or a stock market cycle may have a dominant frequency of about forty one months (cf. Niemira and Klein, 1994, p.268 and p.430). Thus, one way to view a dominant frequency may be as an average cycle length for the data series under investigation.[4]

As the window is moved forward, and paired sequences of observations from adjacent windows are extracted, Ridley (1994) shows that the time series model underlying the system may be viewed as

---

[4] Another way might be to view a dominant frequency (window size for use in the moving window analysis) as that determined by the best fitting model in the spectral regression analysis discussed below.

$$X_t = \sum_{k=1}^{L} X_{(t-k)} B_k + E_t \tag{3}$$

where: B_k is the ordinary least squares coefficient of X lagged k time periods; E_t is the $iid(0,\sigma^2)$ unobservable error term; and t ranges from 1 through N.

Rather than estimate a lag structured model given by equation (3) we may, instead, assume that the variable in the model is generated by a stable stochastic process that depends on different frequencies. Under this assumption Ridley (1994) shows that a complex fast Fourier transform may be applied to each window to estimate the spectral density function from:

$$X_l(\omega) = \sum_{k=1}^{L} X_{(l+L-k)} e^{-i\omega k} \tag{4}$$

where: $l=1,2,...N-L+1$ is the window number and the index of an observation at frequency ω, for $\omega = \pm 0,1,2..., Int(L/2)$ where Int denotes the integer part; e is the exponential base of natural logarithms; and $i = \sqrt{-1}$. In like manner for the system parameters and errors we have:

$$B(\omega) = \sum_{k=1}^{L} B_{(k)} e^{-i\omega k} \tag{5}$$

$$E_l(\omega) = \sum_{k=1}^{L} E_{(l+L-k)} e^{-i\omega k} \tag{6}$$

where each of: X_t and $X_l(\omega)$; B_k and $B(\omega)$; and E_t and $E_l(\omega)$ are complex fourier transform pairs.

Since the spectral densities are symmetric about the zero frequency once the non-negative frequencies are determined the negative frequencies can be obtained directly, implying that only the integer part of $(L/2)+1$ non-negative frequencies are required in the regression model. In transforming the time domain model (eqn.3) to the frequency domain, Ridley (1994) presents the following spectral regression model for estimation:

$$Y_l(\omega) = Y_{l-1}(\omega) + E_l(\omega) \tag{7}$$

for $l=2,3,....N-L+1$ and $\omega = \pm 0,1,2,...,Int(L/2)$ for $-\pi < \omega < \pi$ where $B(\omega)$ is the spectral density function of the impulse response function B_k.

ARIMA Models
The ARIMA approach is an iterative three-stage process of identification, estimation and testing.[5] The general non-seasonal ARIMA (p,d,q) model can be expressed as:

$$\phi(S)\,[(1-S)^d\,X_t - \mu] = \theta(S)\varepsilon_t \tag{8a}$$

---

[5] While autoregressive (AR) models were first introduced by Yule (1926) and moving average (MA) models by Slutzky (1937), it was Wold (1938) who provided the theoretical foundations for combined ARMA processes. Ensuing research considered notions of stationarity and integrated processes, but it was Box and Jenkins (1976) who developed the requisite procedures to effectively understand and use the ARIMA approach to model development and time series forecasting. The Box/Jenkins methodology has been shown by a large number of researchers to produce useful forecasts across diverse data series.

where p refers to the number of autoregressive terms, d is the required degree of differencing for stationarity, q refers to the number of moving average terms, X_t is the series of interest, μ is the mean of the differenced series (often zero), S is the backshift operator and ε_t is a white noise error term. Furthermore, just as consecutive data points might exhibit mixed ARIMA components, so data points separated by a season may show the same properties. The ARIMA notation may be extended to incorporate seasonality and the shorthand notation for the general model becomes ARIMA(p,d,q)*(P,D,Q)[b]   where the lower case refers to the non-seasonal part of the model, the uppercase to the seasonal part and the superscript 'b' refers to the number of periods per season.[6] If we replace the backshift operator, S, by its seasonal counterpart, S^b, the model may be expressed as:

$$\phi\,(S^b)\,[(1 - S^b)^d\,X_t - \mu] = \theta\,(S^b)\varepsilon_t \tag{8b}$$

So, for example, an ARIMA(1,1,0)*(0,1,1)[4] refers to a model operating on quarterly data (indicated by the superscript '4') where the model has a non-seasonal AR(1) component, a non-seasonal difference, no non-seasonal MA component, no seasonal AR component, a seasonal difference and a seasonal MA(1) component.

The ARIMA approach allows for the creation of a broad class of possible data generating processes, and this richness of model representation makes the technique a popular choice amongst forecasters. The basic requirement in the diagnostic checking of a tentatively identified ARIMA model is that: only necessary parameters are included (all of the parameters are statistically significant); parameters satisfy stationarity and invertibility conditions (the series can be represented by a finite order, or convergent autoregressive process); and the parameters included are sufficient (any additional parameters in the model are superfluous in the sense that there are no significant 'patterns' left to exploit).[7] While these diagnostic checks ensure that the tentatively identified model is adequate, they do not ensure that the model is the best ARIMA model for the given data set. In other words there may be another model meeting all of the diagnostic checks that provides a better fit to the data (and, perhaps, better forecasts).

3. Data and Test Results

The analysis was undertaken using quarterly Australian unemployment rate data for the period 1978:2 to 2002:3.[8] There are four reasons for choosing quarterly data over this particular period. First, from February 1978, the Australian Bureau of Statistics (ABS) commenced publishing regular monthly population survey data from which unemployment estimates were culled. Prior to this date, the ABS unemployment rate data were collected on a mid-quarter basis. Focusing on the post 1978:2 period facilitates the application of a consistent data set. Second, the period 1978:2 to 2002:3 encapsulates an era of labour market experience discernibly different from the previous couple of decades during which major structural changes occurred. Granger (1996), among others, recommends choosing a period of

---

[6] The 'b' may be omitted if it is clear from the data series whether it is monthly, quarterly etc.
[7] Refer to Box, Jenkins and Reinsel (1994) for a more complete discussion of model diagnostic checking.
[8] Source: Australian Bureau of Statistics, *The Labour Force, Preliminary*, 6202.0, various issues. Available through EconData Pty. Ltd.

relatively stable attributes when seeking to improve forecast accuracy. McHugh *et al.* (2002) argue along similar lines. Third, extending this analysis sometime in the future by exploring data interdependence with other economic variables, for instance, will likely necessitate the use of quarterly data. Finally, the averaging out of 3-month observations eliminates some of the noise in the data without eliminating the series seasonality or affecting the broad shape of the data over time.

In addition, we divide the unemployment rate series into one for males and one for females. This separation is of interest as over the period there has been a steady improvement in the official unemployment rate for females relative to males.

We next compare the forecast results of a spectral analysis model for our three unemployment rate series with standard ARIMA model forecasts. These comparisons are similar to those carried out by, for example, Rothman (1998) and Stevenson and Peat (2000).

Before applying the spectral and ARIMA tests, let us visually inspect the time series for the unemployment rate for persons, charted in figure 2. The charts for males and females are similar to the charts for persons, so we will confine our observations to the unemployment rate for persons. Seasonally unadjusted and centred four-quarter moving-average data, referred to in figure 2 as 'Smoothed' data, are presented. It can be noted that these data have two prominent cyclical peaks, i.e., two major recessions, occurring on and around 1982:4 and 1990:4

Figure 2 Australian Unemployment Rate for Persons

Seasonally unadjusted ———   Smoothed ━━━
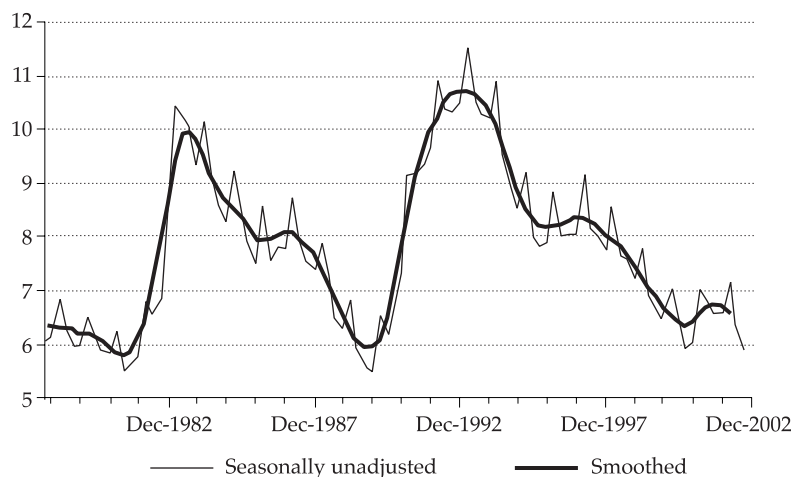
The unemployment rate series used in this study are not stationary. The first differenced series are stationary[9] (see figure 3 for the persons series). The differenced series for seasonally unadjusted persons is seemingly noisy and it is difficult to *visually* discern possible underlying cyclical movements. However, the smoothed series when differenced seems to indicate some

---

[9] As shown by both ADF and PP unit root tests. Spectral methods are applied to stationary series.

underlying cyclical patterns that might be viewed in an approximate fashion as pointing to underlying cycles in the growth rate of the unemployment rate. These smoothed data seem to suggest five cyclical troughs and five cyclical peaks. The number of quarters, between the various peaks, range from 16 to 23 quarters, or four years to 5.75 years.

**Figure 3 Australian Unemployment Rate for Persons First Differenced Data**
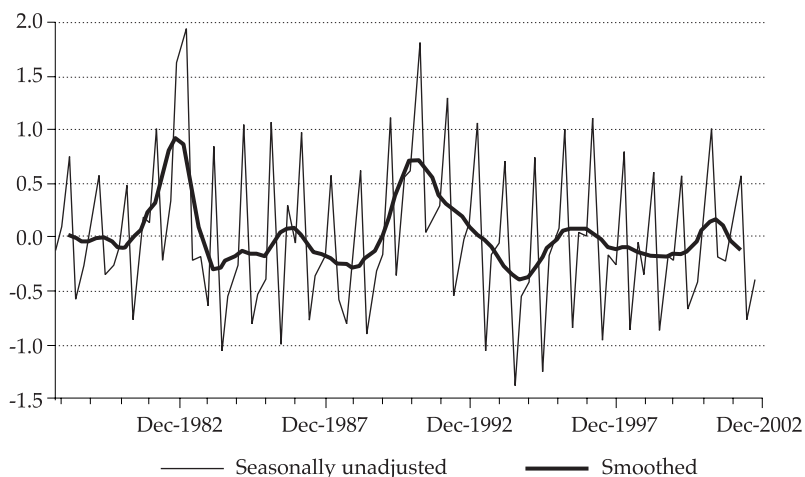
Time Averaging Spectral Estimates

While the literature (above) attests to the acknowledgement amongst economists of the mechanical usefulness of spectral analysis in the decomposition of an economic time series into different frequency components so as to identify primary cycles, long term trends and seasonal influences, there may be some reticence on the part of analysts as to the potential usefulness of spectral analysis for forecasting purposes. Thus, while spectral analysis is useful for a known data set, the question arises: which part of the historically estimated series may provide useful information on the future? Welch (1967) and later Ridley and Mobolurin (1987) provided a potential solution to this dilemma. The time averaging or moving window method seeks to establish a relationship between spectral estimates of the data set at different frequencies, corresponding to current and past values of the data set in the time domain (under some pre-specified moving window structure), that can then be used for forecasting.

Let us consider an intuitive notion of how the moving window spectral regression works. Spectral analysis of the overall data set provides preliminary information on what may be the dominant frequency or primary cycle in the series. So, if the primary cycle for the whole data set is found to be, say, four years then we chose this as the window size to be moved through the full series. Suppose now that instead of the data set as a whole, we chose a window length the size of this dominant cycle (i.e. 16 quarters) for further detailed analysis. If a Fast Fourier Transform is undertaken on this smaller sub-set then the data in this window can be decomposed into its respective frequency components in the same manner that the full set was decomposed

(figure 4 below presents a hypothetical case with a window size of 16). So, for this smaller sub-set of data, we can obtain information on the relationship between the dominant frequency (cycle) and other frequencies in this window frame, where the other frequencies are multiples of this dominant frequency. Within this window, then, the dominant cycle will only complete once. Multiples exist such that the next sub-cycle will complete twice within the window length, the next cycle will complete three times, and so on. Suppose further that the window is now moved forward one time unit and a spectral analysis is similarly undertaken on the new sub-set of data in this second window, and so on through the full data set. Now, if we can establish the relationship between the frequency components of one window with those of the previous window, then we have the basis for examining how the changing alignment of cycles may affect the data set (i.e., be observed as turning points in much the same manner as the hypothetical case in figure 1c). This is essentially what Ridley's (1994) spectral regression is doing. So, for instance, if one hundred quarterly observations are available in the time domain there will be 84 paired observations (windows) available for each spectral regression (since sixteen observations are used up in the first 'roll' through the data set).

Figure 4 Hypothetical Decomposition for First Window
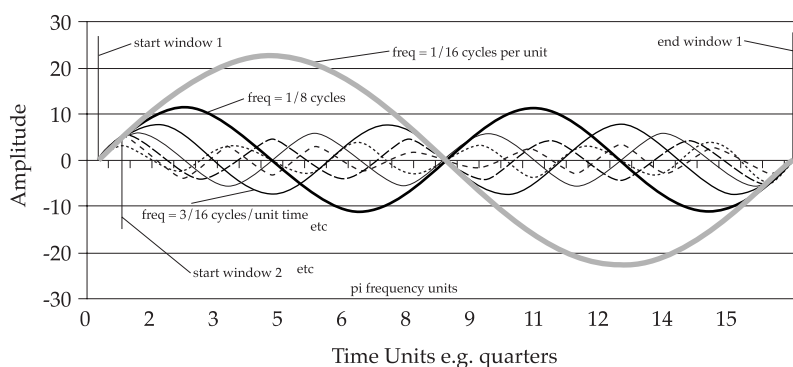
Ridley (1994) argued that a major advantage of the rolling window approach is the ability to pick up the interrelationship between frequencies (cycles of different length), which are difficult to detect in the time domain. If the method predicts an alignment of such cycles in the frequency domain then this may well reflect a trend change in the time domain. An important question, of course, relates to how to determine the size of the window, L, that is rolled through the series. One way around the issue is to optimise on some appropriate goodness of fit parameter, such as the F-statistic for the spectral regression, that is, select the window size that yields the highest R^2 value. This window size is then assumed to yield the dominant cycle. Following this approach, and optimising in the frequency domain, it was found that the most appropriate window size was five years (20 quarters). Interestingly we note that the optimal size of the window, i.e., the dominant cycle, is of a similar length to that observed earlier with reference to the number of periods between the various peaks in the differenced data.[10]

---

[10] It is important to note that our Fourier analysis did *not* work off the smoothed data. The Fourier analysis worked off the raw data after differencing.

Let us now consider the application of this procedure to two-year dynamic forecasts.[11] Figure 5 presents actual data (broken line), in-sample estimates from the moving spectral regression described above (light line) and out-of-sample forecasts for up to eight quarters ahead (heavy line). Forecast accuracy statistics are presented in table 1. While the forecast model building and accuracy statistics in table 1 appear quite sound (and will be discussed in more detail below), the most striking aspect of the analysis is seen in figure 5. Here we note that, while there was some over- and under-shooting by the model during the build phase, most of the downward and upward turns in the series during this model building phase were 'predicted' by the model with very few lag effects. When we examine the out-of-sample forecasting phase we note that, while the forecasts on the level of the series were somewhat poor between 3 and 5 quarters out, virtually all turning points in the series were predicted. In particular, the sharp downward turn in the series some six quarters ahead was predicted by the model, and the forecasts at 1 and 2 and 6 to 8 quarters out can be considered quite sound.
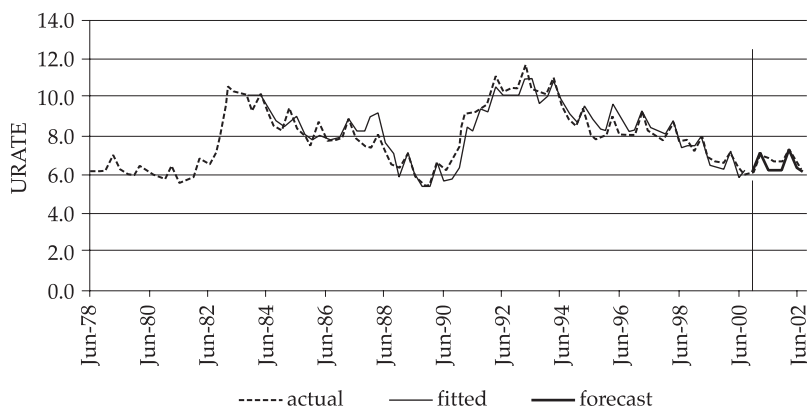
Figure 5 Spectral Forecasts: Persons

------ actual   ——— fitted   ——— forecast

Table 1 Unemployment Rate – Persons: Dynamic Forecasts

Date	Actual	Forecast	%Error	MAPE	RMSE	RMSPE
December-00	6.00	5.97	-0.55	0.55	0.03	0.55
March-01	7.00	6.95	-0.73	0.64	0.04	0.66
June-01	6.80	6.08	-10.52	3.94	0.41	6.28
September-01	6.60	6.07	-8.00	4.95	0.45	6.75
December-01	6.60	6.07	-8.09	5.58	0.46	7.04
March-02	7.10	7.16	0.81	4.78	0.42	6.36
June-02	6.30	6.28	-0.25	4.14	0.39	5.93
September-02	5.90	6.10	3.38	4.04	0.37	5.73

$R^2 = 0.88$ $F(1,68)=488$ Theil's inequality statistic = 0.83. MAPE refers to mean average percent absolute simulation error for the forecast period up to the indicated quarter. RMSE refers to root mean-squared simulation error for the forecast period up to the indicated quarter. RMSPE refers to root mean-squared percent simulation error for the forecast period up to the indicated quarter.

[11] In dynamic forecasting, model updates are based on predicted rather than actual values. Hence the likelihood of forecast error increases with an increasing time horizon. In static, i.e. one-step-ahead forecasts, the model is updated continuously with actual data. We choose to focus on dynamic forecasts on the grounds that achieving accurate dynamic outcomes are more challenging than step-ahead updated forecasts.

A similar analysis was undertaken for the male and female series and the results are shown in figure 6 accompanied by table 2 and figure 7 accompanied by table 3.
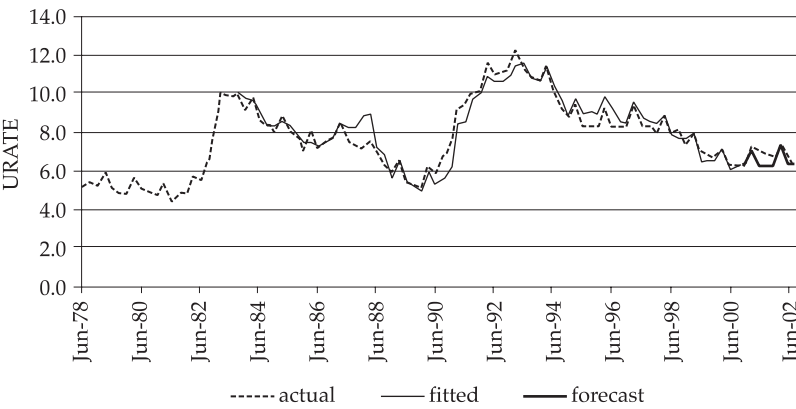
Figure 6 Spectral Forecasts: Males

Table 2 Unemployment Rate Males - Dynamic Forecasts

Date	Actual	Forecast	%Error	MAPE	RMSE	RMSPE
December-00	6.30	6.16	-2.22	2.22	0.14	2.22
March-01	7.30	7.04	-3.53	2.87	0.21	3.04
June-01	7.00	6.12	-12.59	6.11	0.54	7.81
September-01	6.90	6.19	-10.27	7.15	0.58	8.50
December-01	6.80	6.15	-9.54	7.63	0.60	8.71
March-02	7.30	7.15	-2.02	6.69	0.55	7.92
June-02	6.50	6.30	-3.01	6.17	0.51	7.47
September-02	6.10	6.30	3.25	5.80	0.49	7.16

$R^2 = 0.90$ $F(1,68) = 582$ Theil's inequality statistic = 0.94
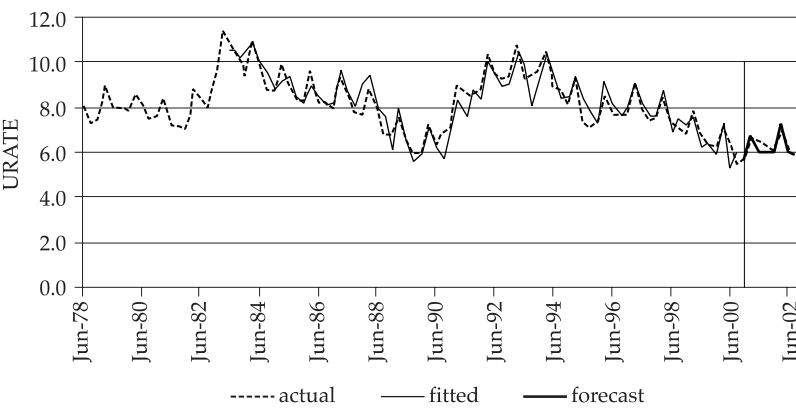
Figure 7 Spectral Forecasts: Females

Table 3 Unemployment Rate – Females - Dynamic Forecasts

Date	Actual	Forecast	%Error	MAPE	RMSE	RMSPE
December-00	5.50	5.58	1.48	1.48	0.08	1.48
March-01	6.70	6.67	-0.45	0.97	0.06	1.01
June-01	6.50	5.95	-8.52	3.48	0.32	5.19
September-01	6.10	5.94	-2.61	3.27	0.29	4.70
December-01	6.30	5.95	-5.55	3.72	0.30	4.89
March-02	6.90	7.16	3.75	3.73	0.30	4.69
June-02	6.10	6.09	-0.12	3.21	0.27	4.36
September-02	5.80	5.80	0.00	2.81	0.26	4.12

$R^2 = 0.83$ $F(1,68) = 323$ Theil's inequality statistic = 0.63

The pattern for male unemployment rates is quite similar to the overall picture. The fitted model again yields reasonably accurate out-of-sample forecasts. Figure 7 and table 3 presents the same analysis for female unemployment rates.

We note that while the model-build statistics (R^2 and F statistics) for females are not quite as strong as the model-build statistics for persons and males, the out-of-sample forecast statistics are (somewhat paradoxically) better, as is clearly evident in figure 7 when compared to figures 5 and 6.

Let us now review and compare some of the summary test statistics for forecasts. The Mean Absolute Percentage Estimation Error (MAPE) and Root Mean Square Percentage Estimation Error (RMSPE) statistics are conventionally used to compare the out-of-sample performance of one model against another. If the statistics are used in isolation (that is, not for comparison purposes), then some subjective interpretation of their quality is required. For example, is the user seeking a model that produces less than five percent error or less than ten percent error etc. in out-of-sample forecasts? Absolute and squared measures are conventionally used to avoid the problem of positive and negative errors cancelling each other, while percentage errors provide a convenient basis for comparison. Comparing tables 1, 2 and 3 shows that the spectral model forecasts are most accurate for females followed by persons followed by males. These differences are encapsulated in the final period RMSPE values of 4.12, 5.73 and 7.16 respectively.

Theil's inequality statistic can be used for further evaluation purposes.[12] The statistic gives a summary measure of how the model's forecasts compare to the forecasts generated from a 'naïve' model that assumes each period's forecast value is the same as the previous period's actual value. The closer Theil's inequality statistic is to zero, the more accurate the forecast. A value of unity suggests the model generates results no better than the naïve model. And values greater than unity indicate model forecasts are less accurate than the naïve case. Ranking of these results favours the forecasts for females over persons over males (0.63, 0.83 and 0.94 respectively).[13]

---

[12] See, Makridakis *et al.* (1991) for further discussion.
[13] Note that the use of Theil's statistic here is a somewhat severe test. This is because the dynamic forecasts from the model do *not* utilise realised data from the forecast period, whereas the 'naïve' alternative *does* utilise realised values from the forecast period, even though these are lagged realised values.

ARIMA Estimates

We next develop ARIMA models for the purpose of comparison with our
spectral model results. This approach is similar to comparisons with linear
autoregressive models made by Stevenson and Peat (2000). The ARIMA
models detailed below represent models that satisfy all of the diagnostic
checks as outlined earlier.[14]

**Figure 8 ARIMA Forecasts for Persons: Unemployment Rate ARIMA
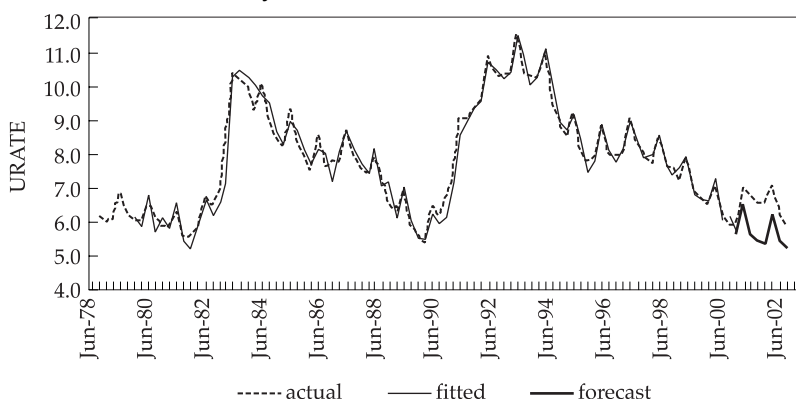(1,1,0)*(0,0,1) Model, Dynamic Forecasts**

Table 4 Unemployment Rate – Persons – ARIMA(1,1,0)*(0,1,1) Dynamic Forecasts

Date	Actual	Forecast	%Error	MAPE	RMSE	RMSPE
Dec-00	6.00	5.67	-5.47	5.47	0.33	5.47
Mar-01	7.00	6.48	-7.47	6.08	0.44	6.72
Jun-01	6.80	5.64	-17.00	9.84	0.76	11.46
Sep-01	6.60	5.44	-17.59	12.00	0.88	13.26
Dec-01	6.60	5.35	-19.00	13.40	0.96	14.59
Mar-02	7.10	6.23	-12.24	12.42	0.95	14.19
Jun-02	6.30	5.44	-13.59	13.94	0.94	14.11
Sep-02	5.90	5.27	-10.73	14.36	0.90	13.82
				Model Parameters*		
		Variable	Coef.	Std.Error	t-stat	
		p[1]	0.5896	0.0878	6.7162	
		Q[1]	0.9283	0.0419	22.1427	

$R^2 = 0.96$ F(1,88) = 2090 Theil's inequality statistic = 1.3
* Lower case p(1) indicates one non-seasonal AR component (similarly lowercase q(1) would
indicate one non-seasonal MA component). Uppercase Q(1) indicates one seasonal MA component
(similarly P(1) indicates one seasonal AR component. Note here the seasonal is quarterly).

The model presented in figure 8 along with table 4 is an ARIMA(1,1,0)*(0,1,1)
seasonal model with the coefficient estimates as shown in table 4. Figure 8
presents dynamic (i.e., 8-steps ahead) forecasts. In a direct comparison with
the spectral regression method (tables 1 and 5) we can see that the in-
sample model fit statistics are better for the ARIMA model, while the out-

---

[14] However there may be alternative ARIMA model specifications that may either
under- or out-perform our models. Some care should be taken, therefore, in the
interpretation of the forecasting comparisons. Again we note that a full discussion
of diagnostic checks can be obtained from Box, Jenkins and Reinsel (1994).

of-sample forecast accuracy statistics are better for the spectral regression model. Figures 5 and 8 show that in-sample turning point prediction is quite comparable, although the out-of-sample turning point prediction again is better for the spectral method. It can be seen that the spectral model is closer in its identification of a possible out-of-sample turning point, and on this basis out-performs this particular ARIMA model. Figure 5 shows that the spectral model forecasts a possible upturn in the series, as well as identifying seasonal fluctuations at about the time that these actually occur. The ARIMA model does not pick the potential out-of-sample turning point, maintaining the most recent linear downward trend, although it does very well on seasonal fluctuations.[15] The Theil inequality statistic similarly points to the relative accuracy of the spectral forecasts over the ARIMA forecasts.

Figure 8 ARIMA Forecasts for Males: Males - Unemployment Rate ARIMA (1,1,0)*(0,0,1) Model, Dynamic Forecasts
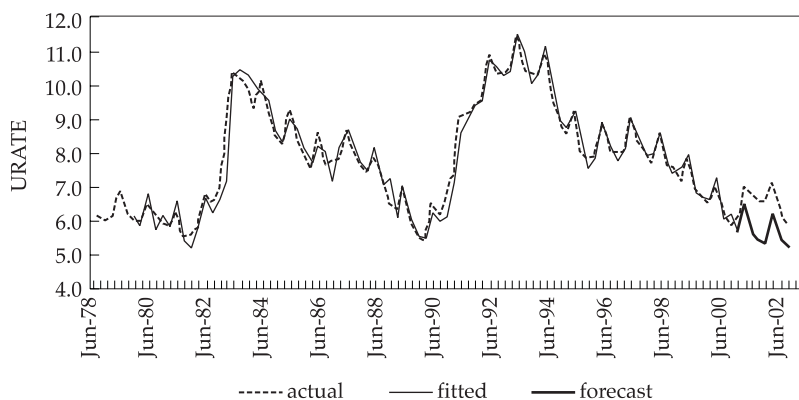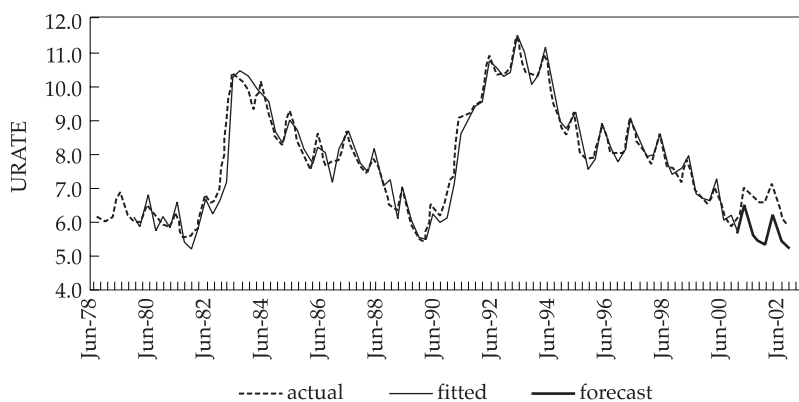

Figure 9 ARIMA Forecasts for Males: Males - Unemployment Rate ARIMA (1,1,0)*(0,0,1) Model, Dynamic Forecasts

---

[15] It should be borne in mind that no intervention analysis (transfer function) was undertaken as part of this modelling procedure. This may well have improved model estimation (and performance) and that will be the subject of later research.

**Table 5 Unemployment Rate – Males – ARIMA(1,1,0)*(0,1,1) Dynamic
Forecasts**

Date	Actual	Forecast	% Error	MAPE	RMSE	RMSPE
Dec-00	6.30	6.00	-4.75	4.75	0.30	4.86
Mar-01	7.30	6.67	-8.59	6.34	0.46	6.99
Jun-01	7.00	5.95	-14.96	9.40	0.74	11.01
Sep-01	6.90	5.89	-14.71	10.83	0.84	12.50
Dec-01	6.80	5.76	-15.37	11.86	0.91	13.57
Mar-02	7.30	6.48	-11.30	11.09	0.90	13.26
Jun-02	6.50	5.79	-10.95	12.24	0.89	13.14
Sep-02	6.10	5.74	-5.84	12.14	0.85	12.76
			Model Parameters			
		Variable	Coef.	Std.Error	t-stat	
		p[1]	0.70	0.08	8.99	
		Q[1]	0.94	0.04	23.70	

$R^2 = 0.98$ F(1,88) = 3404 Theil's inequality statistic = 1.5
* Lower case p(1) indicates one non-seasonal AR component (similarly lowercase q(1) would
indicate one non-seasonal MA component). Uppercase Q(1) indicates one seasonal MA component
(similarly P(1) indicates one seasonal AR component. Note here the seasonal is quarterly).
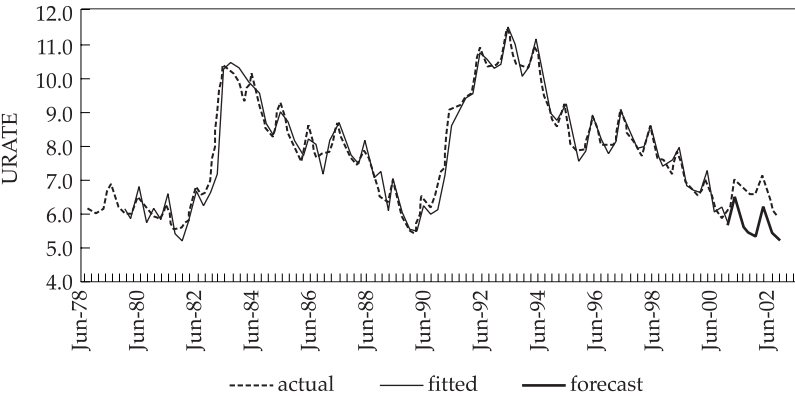
Males

The ARIMA model estimates for the male unemployment rate are shown
in figure 9 and table 5. Once again an ARIMA(1,1,0)*(0,1,1) model satisfied
all of the diagnostics as presented earlier. Again the in-sample model-build
statistics are much better for the ARIMA model while the out-of-sample
forecast-accuracy statistics are not as good as those presented earlier for
the spectral regression for the males series (figure 6 and table 2). The ARIMA
model again failed to pick the possible upturn in the series.

Females

The ARIMA model estimates for the female unemployment rate are shown
in figure 10 and table 6. Once again an ARIMA (1,1,0)*(0,1,1) model satisfied
all of the diagnostics. As with the data sets for persons and males, the in-
sample model-build statistics are considerably superior for the ARIMA
model than the spectral model. However, the out-of-sample forecast
accuracy statistics for the ARIMA model are not as good as those for the
spectral model. This is confirmed by a visual inspection of the charts as
well as the various test statistics.

**Figure 10 ARIMA Forecasts for Females: Females - Unemployment
Rate ARIMA (1,1,0)*(0,0,1) Model, Dynamic Forecasts**

**Table 6 Unemployment Rate – Females – ARIMA (1,1,0)*(0,1,1)
Dynamic Forecasts**

Date	Actual	Forecast	%Error	MAPE	RMSE	RMSPE
Dec-00	5.5	5.317	-3.33	3.33	0.18	3.08
Mar-01	6.7	6.407	-4.37	3.55	0.40	6.26
Jun-01	6.5	5.506	-15.29	7.54	0.70	10.65
Sep-01	6.1	5.148	-15.61	9.93	0.80	12.20
Dec-01	6.3	5.129	-18.59	11.41	0.88	13.48
Mar-02	6.9	6.238	-9.59	10.28	0.87	13.07
Jun-02	6.1	5.340	-12.46	11.74	0.86	12.98
Sep-02	5.8	4.983	-14.09	12.57	0.83	12.77
				Model Parameters		
		Variable	Coef.	Std.Error	t-stat	
		p[1]	0.1915	0.1062	1.8033	
		Q[1]	0.9400	0.0302	31.1405	

$R^2 = 0.91$ $F(1,88) = 873$ Theil's inequality statistic = 1.3

What forces, then, might be responsible for these seemingly anomalous results of spectral analysis performing relatively well out-of-sample, but relatively poorly in-sample? We suggest that two influences may be operating. First, as a general observation, our results demonstrate that the accuracy of in-sample predictions should not be overplayed. The illusion of accuracy associated with successful in-sample predictions is akin to wisdom gained *after the fact*. Successfully modelling *before the fact* is much more challenging. Being able to account for or rationalise the past may not necessarily be a good guide to modelling the future. If a model captures, in a crude way, underlying cyclical forces at play, and if these myriad forces are so complex to only ever be approximately distilled in a mathematical form, it may turn out that a seemingly less accurate model of in-sample outcomes nevertheless proves to be a more accurate model of out-of-sample outcomes. Arguably, for the timeframe and data employed in this analysis, spectral analysis has done exactly this.

Our second observation is more specific to the mathematics of spectral and ARIMA models. The interacting cycles of a spectral model are capable of generating irregular overall changes in a variable and projecting these changes into the future. For example, if two or more cycles of different frequency and amplitude happen to coincide when they are (say) peaking, then the peak overall value of the variable will appear to be abnormally high. At first glance the overall pattern of change in the variable may appear to be so irregular as to be seemingly unpredictable. The ARIMA model chosen to model unemployment, on the other hand, displays a regular future cyclical path. Arguably it is the apparent irregularity generated by the moving window spectral model that has generated forecasts for the period under review that better mimic the apparent irregularities of the realised changes in the unemployment rate.

Perhaps there is a general lesson in these results: that impressive modelling of the past does not guarantee accurate forecasts of the future. After all, how many models of past relations have survived *in their detail* into the future? It is no secret that the answer is few, if any, have survived. Perhaps,

then, the spectral approach has picked up in an approximate fashion
unfolding cyclical forces to which the ARIMA model is not as well attuned.
That, in any case, is the view we take in this paper for these data.

4. Concluding Thoughts

Both the ARIMA and spectral regression modelling processes are designed
to mathematically capture, from the past, forces, patterns and regularities
that may carry through into the future. In our univariate modelling, the
complicated actions and reactions that generate unemployment rate
outcomes are largely unknown in terms of their detailed machinations.
Perhaps they are unknowable. Certainly they appear to be unknowable in
their detail. The question we have posed in this paper is whether spectral
analysis – thus far ignored in the univariate forecasting research on
unemployment – is worth consideration as a forecasting model.

We have answered that question in the affirmative. Spectral modelling for
the series and the time frame considered, employing the moving window
spectral regression method refined by Ridley and Mobolurin (1987) and
Ridley (1994), has performed more than passably. Specifically, we have
found that our spectral models have produced more accurate out-of-sample
forecasts than both (i) an estimated ARIMA model and (ii) 'naïve' forecasts
based on a set of one-period-prior realised values. The picture, though, is a
little more complex than spectral model results being unambiguously
superior to all its rivals. The ARIMA model produced superior in-sample
model-build statistics. We have suggested that this may indicate that, both
as a general observation and an observation specific to this study, generating
strong in-sample predictions is not necessarily a good guide to how accurate
out-of-sample forecasts will be. We have argued that despite the spectral
models explaining less in-sample variability than the ARIMA model, the
spectral model may have nevertheless come closer to capturing the
underlying machinations of the past unfolding into the future.

Finally, we note that this study is a preliminary study; not only in the sense
that as events unfold and techniques develop, new methodologies will
emerge, but also in the sense that further applications of spectral analysis
are worth exploring with regards to unemployment rate series for other
economies as well as searching out long-term regularities in cyclical
behaviour. In addition, spectral results need to be compared with model
forecasts other than ARIMA, given that the results of this paper suggest
that spectral analysis is indeed a worthy contender.

References

Acemoglu, D. and Scott, A. (1994), 'Asymmetries in the Cyclical Behaviour
of UK Labour Markets', *Economic Journal,* 104, 1303-1323.

Ball, L. (1997), 'Disinflation and the NAIRU', in Romer, C. and Romer, D.
(eds.), *Reducing Inflation,* 167-185, University of Chicago Press.

Blanchard, O. (1991), 'Wage Bargaining and Unemployment Persistence',
Journal of Money, Credit and Banking, 23, 277-292.

Bloomfield, P. (1976), *Fourier Analysis of Time Series: An Introduction,* Wiley,
New York.

Bodman, P. (1998), 'Asymmetry and Duration Dependence in Australian GDP and Unemployment', *Economic Record*, 74, 399-412.

Bodman, P. (2001), 'Steepness and Deepness in the Australian Macroeconomy', *Applied Economics*, 33, 375-382.

Borland, J. (1997), 'Unemployment in Australia - Prospects and Policies: An Overview', *The Australian Economic Review*, 30, 391-404.

Box, G.E.P. and Jenkins, G.M. (1976), *Time Series Analysis: Forecasting and Control*, Revised Ed., San Francisco, Holden Day.

Box, G.E.P., Jenkins, G.M. and Reinsel, G. (1994), *Time Series Analysis: Forecasting and Control*, 3rd Ed., San Francisco, Prentice-Hall, Englewood-Cliffs, NJ.

Burtless, G. (2002), 'Can Supply-Side Policies Reduce Unemployment? Lessons from North America', *Australian Journal of Labour Economics*, 5, 115-142.

Chapman, B. and Kenyon, P. (2002), 'Creating Jobs – the Role of Government: Introduction', *Australian Journal of Labour Economics*, 5, 101-110.

Dawkins, P. (2002), 'The 'Five Economists' Plan: The Original Idea and Further Development', *Australian Journal of Labour Economics*, 5, 203-230.

Debelle, G. and Vickery, J. (1998), 'The Macroeconomics of Australian Unemployment', in Debelle, G. and Borland, J. (eds.), *Unemployment and the Australian Labour Market*, 235-265, Reserve Bank of Australia, Sydney.

Enders, W. (1995), *Applied Econometric Times Series*, John Wiley & Sons, New York.

Engle, R.F. (1974), 'Band Spectrum Regression', *International Economic Review*, 15, 1-11.

Forslund, A. and Krueger, A. (1997), 'An Evaluation of Swedish Active Labour Market Policy: New and Received Wisdom', in Freeman, R., Topel, R. and Swedenborg, B. (eds.), *The Welfare State in Transition*, 267-298, University of Chicago Press, Chicago.

Granger, C.W.J. (1996), 'Can We Improve the Perceived Quality of Economic Forecasts', *Journal of Applied Econometrics*, 11, 455-473.

Granger, C.W.J. (1969), 'Investigating Causal Relations by Econometric Models and Cross-Spectral Methods', *Econometrica*, 37, 424-438.

Gregory, R.G. (2002), 'Round Table Discussions Its Full-Time Jobs that Matter', *Australian Journal of Labour Economics*, 5, 279-288.

Hamilton, James D. (1994), *Time Series Analysis*, Princeton University Press, New Jersey.

Hannan, E.J. (1963), 'Regression for Time Series', in Rosenblatt, M. (ed.), *Time Series Analysis*, New York, John Wiley.

Hannan, E.J. (1965), 'The Estimation of Relationships Involving Distributed Lags', *Econometrica*, 33, 206-224.

Hannan, E.J. (1967), 'The Estimation of a Lagged Regression Relation', *Biometrika*, LIV, 409-418.

Harvey, A.C. (1978), 'Linear Regression in the Frequency Domain', *International Economic Review*, 19, 507-512.

Iksoon, E., Cauley, J. and Sandler, T. (1987), 'Cycles and Substitutions in Terrorist Activities: A Spectral Approach', *Kylklos*, 40, 238-255.

Jackman, R. (1998), 'European Unemployment: Why is it So High and What Should be Done', in Debelle, G. and Borland, J. (eds.), *Unemployment and the Australian Labour Market*, 39-63, Reserve Bank of Australia, Sydney.

Katz, L.F. (1998), 'Reflections on US Labour Market Performance', in Debelle, G. and Borland, J. (eds.), *Unemployment and the Australian Labour Market*, 8-33. Reserve Bank of Australia, Sydney.

Layard, R., Nickell, S. and Jackmann, R. (1991), *Unemployment, Macroeconomic Performance and the Labour Market*, Oxford University Press, Oxford.

Levy, Daniel (2002), 'Cointegration in Frequency Domain', *Journal of Time Series Analysis*, 23, 333-339.

Lindbeck, A. (1995), 'Harzardous Welfare-state Dynamics', *American Economic Review*, 85, 9-15.

Makridakis, S.G., Wheelwright, S.C. and McGee, V.E. (1991), *Forecasting: Methods and Applications*, 2nd ed., New York, Wiley.

McHugh, Z., Bardsen, G. and Hurn, S. (2002), 'Modelling Unemployment in Australia', Working Paper No. 10/2002, Department of Economics, Norwegian University of Science and Technology.

McNees, S.K. (1982), 'The Role of Macroeconometric Models in Forecasting and Policy Analysis in the United States', *Journal of Forecasting*, 1, 37-48.

Neftci, S. (1984), 'Are Economic Time Series Asymmetric Over the Business Cycle?' *Journal of Political Economy*, 92, 307-328.

Nerlove, M. (1964), 'Spectral Analysis of Seasonal Adjustment Procedures', *Econometrica*, 32, 241-286.

Niemira, M.P. and Klein, P.A. (1994), *Forecasting Financial and Economic Cycles*, John Wiley and Sons, NY.

OECD (1994), *Employment Outlook*, OECD, Paris.

Parker, R.E. and Rothman, P. (1997), 'The Current Depth-of-recession and Unemployment-rate Forecasts', *Studies in Nonlinear Dynamics and Econometrics*, 2, 151-158.

Perry, L.J. and Wilson, P.J. (2001), 'Strikes and the Accord: An International Perspective', *Australian Journal of Labour Economics*, 4, 232-245.

Ridley, A.D. (1994) 'The Global Univariate Moving Spectral Window Method', *Supercomputer Computations Research Institute*, Florida State University, February,

Ridley, A.D. and Mobolurin, A. (1987), 'Application of a Spectral Analysis Algorithm (MWS) to Time Series Forecasting of Automobile Demand', *Proceedings, The Twenty-Fourth Annual Symposium of the Washington Operations Research/Management Science Council*, Nov.

Rothman, P. (1998), 'Forecasting Asymmetric Unemployment Rates', *Review of Economics and Statistics*, 80, 164-168.

Siebert, H. (1997), 'Labour Market Rigidities: At the Root of Unemployment in Europe', *Journal of Economic Perspectives*, 11, 37-54.

Skalin, J. and Terasvirta, T. (2002), 'Modeling Asymmetries and Moving Equilibria in Unemployment Rates', *Macroeconomic Dynamics*, 6, 202-241.

Slutzky, E. (1937), 'The Summation of Random Causes as the Source of Cyclic Processes', *Econometrica*, 5, 105-146.

Stevenson, M. and Peat, M. (2000), 'Forecasting Australian Unemployment Rates', *Australian Journal of Labour Economics*, 4, 41-55.

Valentine, T. (1993), 'The Sources of Unemployment: A Simple Econometric Analysis', *Economic Papers*, 12, 1-20.

Van Dijk, D., Franses, P.H. and Paap, R. (2002), 'A Nonlinear Long Memory Model with an Application to US Unemployment', *Journal of Econometrics*, 110, 135-165.

Verbrugge, R. (1997), 'Investigating Cyclical Asymmetries', *Studies in Nonlinear Dynamics and Econometrics*, 2, 15-22.

Welch, P. (1967), 'The Use of Fast Fourier Transform for Estimation of Power Spectra: A Method of Time Averaging over Short, Modified Periodograms', *IEEE Transactions in Audio Electroacoustics*, AU-15, 70-3

Wilson, P. J., Okunev, J., Ellis, C. and Higgins, D.M (2000), 'Comparing Univariate Forecasting Techniques in Property Markets', *Journal of Real Estate Portfolio Management*, 6, 283-306.

Wilson, W. (1996), *When Work Disappears: The World of the New Urban Poor.* Knopf, New York.

Wold, H. (1938), *A Study in the Analysis of Stationary Time Series*, Stockholm, Almquist and Wiksell.