

The GSBGM Working Paper Series

WORKING PAPER 1/98

**Selection and Estimation of Trigonometric Component
Models for Seasonal Time Series**

John Haywood and Granville Tunnicliffe Wilson

© 1998 John Haywood and Granville Tunnicliffe Wilson.

All rights reserved. Apart from any fair dealing for the purpose of private study, research, criticism or review, as permitted under the Copyright Act, no part of this document may be reproduced by any means, electronic, mechanical electrostatic photocopying or otherwise, or stored in a retrieval system without the prior permission of the authors.

ISSN 1173-4523

ISBN 0-475-11527-9

The GSBGM Working Paper Series 1/98 January 1998.

John Haywood* and Granville Tunnicliffe Wilson 'Selection and Estimation of Trigonometric Component Models for Seasonal Time Series'

Please address inquiries concerning this paper direct to the first named author*

Other queries to:

Angela Dolan

Manager Administration Services

The Graduate School of Business and Government Management

The Victoria University of Wellington

PO Box 600

Wellington

New Zealand

Telephone: 64-4-471 5366

Facsimile: 64-4-496 5435

E-mail: angela.dolan@vuw.ac.nz

John Haywood (email John.Haywood@vuw.ac.nz)

School of Economics and Finance

Victoria University of Wellington

PO Box 600

Wellington

New Zealand

Printed by The Victoria University of Wellington Printers.

Selection and Estimation of Trigonometric Component Models for Seasonal Time Series

John Haywood* and Granville Tunnicliffe Wilson†

Abstract

We present a method for investigating the evolution of trend and seasonality in an observed time series. A general model is fitted to a residual spectrum, using trigonometric components to represent the seasonality. We show graphically how well the fitted spectrum captures the evidence for evolving seasonality associated with the different seasonal frequencies. After fitting a seasonal IMA model, the method requires only ordinary least squares estimation. A submodel which adequately fits the data can then be conveniently selected. We apply the method to two time series and discuss the implications for time series forecasting.

Keywords

Structural model; seasonal time series; variance components; trigonometric components; frequency domain estimation.

*School of Economics and Finance, Victoria University, PO Box 600, Wellington, NZ.

E-mail: John.Haywood@vuw.ac.nz

†Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, UK.

E-mail: g.tunnicliffe-wilson@lancaster.ac.uk

1 Introduction

Practitioners of forecasting and seasonal adjustment commonly deal with time series exhibiting trend and seasonality which evolve through time. Among the models and methods used to represent and forecast these series are (a) the seasonal ARIMA models of Box and Jenkins (1976), particularly that known as the *airline* model, (b) the Holt-Winters seasonal forecasting procedure which in its additive form is equivalent to a particular ARIMA model, as shown in McKenzie (1976), and (c) state space models with seasonal components as advocated by Harrison and Stevens (1976) and more recently by Harvey (1990). Some of the earliest models for evolving seasonality were formulated by Hannan (1964) and a similar approach is found in the dynamic harmonic regression models presented by Ng and Young (1990). We shall confine ourselves to linear prediction models in this paper so will not consider the *mixture process* aspects of Harrison and Stevens work, designed to model outliers and sudden changes in trend and seasonality, although the conclusions of this paper may well be relevant to current developments of these by Bruce and Jurke (1996).

When these models are fitted to a time series or, in the case of the Holt-Winters procedure, *tuned* for optimal forecasting performance, the parameters will reflect the extent to which the level, trend and seasonality are evolving through time. The aim of this paper is to investigate the evidence to be found in an observed time series for such evolving features, and to use a flexible spectral component model to represent them. This model may then be applied to forecasting or seasonal adjustment. We shall confine ourselves to considering seasonality with a period of 12 because this is so important in practice, but the methods can be applied equally well to other periods.

Our approach is to fit a very general model to the observed series; a model with 14 parameters of which 12 are used to characterise the evolution of the seasonality including the level, one for evolution of the trend and one for white noise error. These parameters are the coefficients of linear components of the spectrum of the series (strictly the pseudo-spectrum; for convenience we shall use just the term spectrum). The seasonal coefficients in particular are associated with what we shall call *trigonometric components* of the spectrum located at the seasonal frequencies. These trigonometric components *may* be associated with independent unobserved components of the time series *if* their coefficients are positive. (Where necessary to avoid confusion between the components of the spectrum and the series, we shall emphasise when we are referring to series components). For the unobserved components interpretation we present both a simple ARIMA representation and a convenient state space representation of the series components.

Our general model is equivalent to a seasonal IMA model with 13 moving average parameters and one prediction error variance parameter. This encompasses many of the other models referred to above and Newbold (1988) uses it for comparing Box-Jenkins, Holt-Winters and structural models. Our starting point is to fit this model to the data, a quick and reliable procedure using standard packages, and then to use it as a basis for fitting the linear spectrum model to the sample spectrum of the observed series. We propose our own modified method of doing this, which we believe compares well with other methods in that it uses only least squares estimation. This approach also provides a valuable graphical means for displaying how well the fitted spectrum captures the evidence for evolving seasonality.

Regression methods can then be used to test which coefficients are significant and therefore, within the series, which components of evolving seasonality are supported by the data. We illustrate the method for two time series and discuss the implications for the standard models mentioned above.

2 Initial investigation of varying seasonality

We shall use two examples of US macroeconomic series for illustration. The first of these is the inventory of unfilled orders for newspapers and magazines from January 1964 to March 1989 inclusive. Figure 1a shows a graph of the first 25 years of this series after transformation by taking natural logarithms and multiplying by 1000. In this series there appear to be clear changes in both trend and seasonality. To confirm this we can fit a standard regression on fixed trend and seasonality and analyse the departures from this model, ie the errors n_t . For data $\{y_t; t = 1, 2, \dots, n\}$ the model may be expressed

$$y_t = c + bt + \sum_{j=1}^5 (A_j S_{j,t} + B_j C_{j,t}) + B_6 C_{6,t} + n_t \quad (1)$$

where the regressors apart from the constant and trend are the cycles with fundamental period 12 defined by

$$S_{j,t} = \sin 2\pi jt/12 \quad ; \quad C_{j,t} = \cos 2\pi jt/12. \quad (2)$$

An equivalent form would be to replace the constant and trigonometric terms of the regression by a seasonal factor, ie 12 monthly indicator variables.

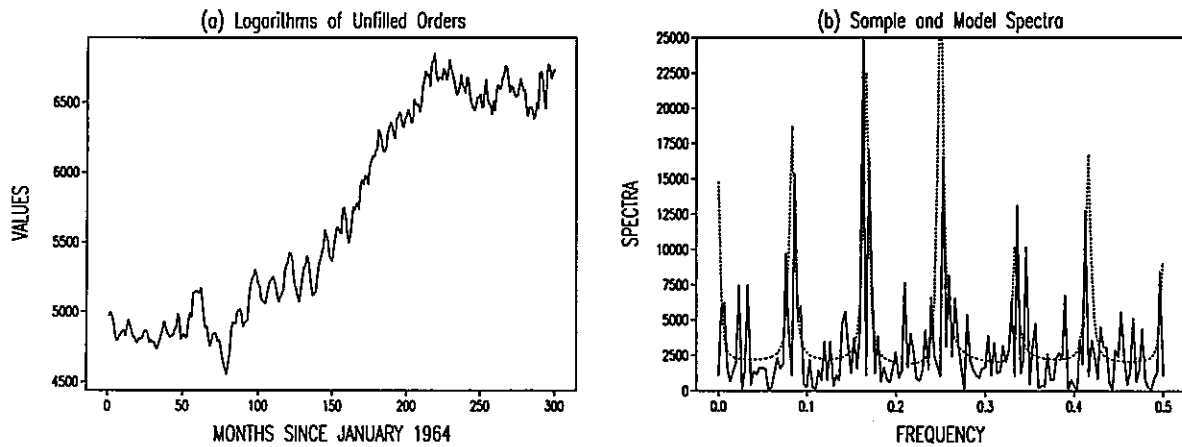


Figure 1: (a) Monthly values of the unfilled orders for newspapers, (b) Sample and model spectrum of this series.

Evidence for changing seasonality would be revealed by irregular peaks close to the seasonal frequencies $F_j = 1/12, 2/12, \dots, 6/12$ in the sample spectrum of the errors n_t from fitting this model. Unfortunately, compared with the low frequency components of these errors arising from the large variations about the fixed trend, these seasonal peaks are typically small and not easily inspected when plotted without modification of the spectrum.

The modification we propose is achieved by first approximating the whole of the correlation in the errors by a simple autoregressive structure

$$n_t = \rho n_{t-1} + e_t \quad (3)$$

whose parameter ρ we determine as the first lag sample autocorrelation of the errors n_t . We then calculate residuals from this autoregression as the *generalized difference* $e_t = n_t - \rho n_{t-1}$ and inspect the sample spectrum of these residuals for the evidence of changing seasonality. This is equivalent to scaling down the low frequencies in the spectrum of n_t . What we have done is to carry out one

step of the Cochrane-Orcutt procedure for fitting the regression (1) with autoregressive errors (3). We have checked for our examples that nothing is lost by using just the one step.

The irregular line in Figure 1b shows the sample spectrum of the residuals e_t , derived from the series in Figure 1a using $\rho = 0.9691$, defined as

$$W^*(f) = \left| \sum_{t=1}^n e_t \exp(2\pi i f t) \right|^2 / n \quad (4)$$

and graphed against frequency f for $0 \leq f \leq 0.5$. The sample spectrum clearly reveals irregular peaks around the seasonal frequencies as evidence of changing seasonality.

If however there were no indications of *excess power*, ie peaks, in this sample spectrum in the neighbourhood of the seasonal frequencies then the fixed seasonal regressors could be considered adequate for the data. If the fixed seasonal terms were *not* first removed by these seasonal regressors then they would appear in this spectrum as large discrete components precisely at the seasonal frequencies. The peaks which evidence changing seasonality would appear as relatively small sidebands to them. It is partly for this reason that we first remove *fixed* seasonal components; we consider that they should not contribute to modelling the *variations* in seasonality.

We shall describe in section 4 how we fit a component spectrum model to this residual sample spectrum. Particular terms in this model aim to explain any excess variation at the seasonal frequencies. We will also display the fitted spectrum values against this residual spectrum as a visual check. The dotted line superimposed on the sample spectrum in Figure 1b is from a preliminary model whose derivation we describe in section 4.

The coefficients and fitted values of our model in fact describe the extent to which the components of varying seasonality are present in the errors n_t of the original series, not in the residual series e_t . It is however a simple matter to scale the fitted values in order to display them against the spectrum of the residuals and thereby gain the advantage of a clearer comparison between data and model, avoiding the dominating effect of the low frequencies.

3 The general model and its components

Our stated aim is to model the variations in the trend and seasonality of the series about the *fixed* component of trend and seasonality as expressed in (1). None of the models which we mentioned in section 1 do in fact explicitly invoke such a fixed component because they are all expressed in terms of the changes from the past to the future. A fixed component of trend and seasonality is, strictly, not even uniquely defined for these models although they do all encompass model (1) as a special or limiting case. All these models do however contain, in their parameters, a statistical description of the variations in trend and seasonality and a convenient general model by which they may all be represented is the seasonal IMA process:

$$\nabla \nabla_{12} y_t = w_t = (1 - \theta_1 B - \theta_2 B^2 - \dots - \theta_{13} B^{13}) e_t = \theta(B) e_t \quad (5)$$

where e_t is white noise and we are using the notation of Box and Jenkins (1976). An equivalent description is that the differenced series $w_t = \nabla \nabla_{12} y_t$ is autocorrelated to a maximum lag of 13. The model may be represented by its *spectrum* which is well defined at all points except frequency zero and the seasonal frequencies F_j defined above, by setting $B = \exp(i\omega)$, where $\omega = 2\pi f$, in

$$\left| \frac{\theta(B)}{\nabla \nabla_{12}} \right|^2 \sigma_e^2. \quad (6)$$

This may be expressed as the rational cosine function

$$P(f) = \frac{\gamma(0) + 2\gamma(1) \cos \omega + \dots + 2\gamma(13) \cos 13\omega}{4(1 - \cos \omega)(1 - \cos 12\omega)} \quad (7)$$

where $\gamma(k)$ are the autocovariances of the differenced series w_t and the numerator is the spectrum of this series. The denominator is the *gain* of the differencing operators $\nabla\nabla_{12}$.

Note that 14 free coefficients appear in both the IMA model (as 13 moving average coefficients and a residual variance) and in the spectrum (as 14 autocovariances). Our next step is to express this spectrum as the unique linear combination of 14 components:

$$P(f) = \sum_{k=0}^{13} V_k P_k(f) \quad (8)$$

where the free parameters are now the coefficients V_k . Each of the components $P_k(f)$ corresponds to the spectrum of a low order process. We describe these components here in simple terms, giving the technical definitions in section 5.

The first component $P_0(f)$ is the only one which does not have a peak. It is the spectrum of white noise with unit variance, so takes the value one over the whole frequency range. The remaining components have peaks which correspond to the zeros of the denominator of $P(f)$ which are double zeros at frequencies $0, F_1, F_2, \dots, F_5$ and a single zero at frequency $F_6 = 0.5$.

Both $P_1(f)$ and $P_2(f)$ have an infinite peak at frequency zero and correspond respectively to the spectrum of a random walk (whose first difference is white noise) and an integrated random walk (whose first difference is a random walk). These model respectively the variation in level and trend of the series.

The remaining eleven components $P_3(f)$ to $P_{13}(f)$ model the variations in seasonality of the series. The first ten of these are considered as five pairs which have peaks at the seasonal frequencies F_1 to F_5 . Figure 2 illustrates these selectively in two frames, the first (second) of which shows the first (second) member of each pair $P_3(f), P_4(f)$ and $P_{11}(f), P_{12}(f)$. Rather than truncate the infinite peaks in the plots they are rounded off so that each corresponds to the spectrum of a stationary ARMA(2,1) process defined in section 5. The members of each pair look very similar in the centre of the peak but are skewed differently, reflecting an important property that the first (second) of each pair falls to zero at the upper (lower) limit of the frequency range. As we shall see, both members of each pair are needed if we are adequately to represent some typical model spectra. In practice we shall fit the models by replacing each pair by its sum and difference, to avoid the high collinearity which otherwise occurs.

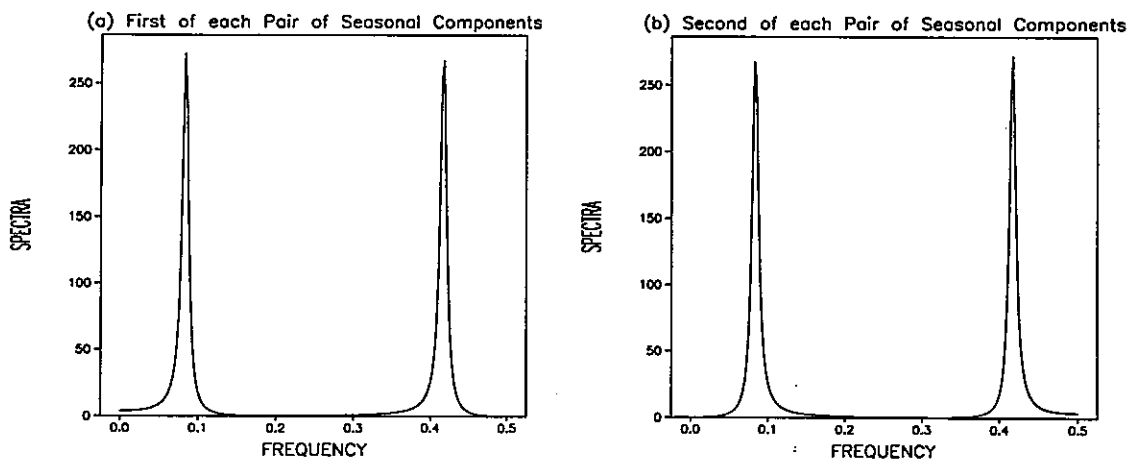


Figure 2: Seasonal spectrum component pairs (a) First components of pairs with peaks at $F_1 = 1/12$ and $F_5 = 5/12$, (b) Second components of the same pairs.

The last, single, seasonal component $P_{13}(f)$ has a peak at frequency 0.5 which mirrors the peak at frequency zero in the random walk component $P_1(f)$.

We end this section by illustrating how the seasonal part of some simple models is represented in terms of these components. In Figure 3a we show the seasonal part of the spectrum of a typical Box-Jenkins seasonal model, the 'airline' model with parameters $\theta = 0.4$ and $\Theta = 0.55$. The peaks at frequency zero associated with the components $P_1(f)$ and $P_2(f)$ have been subtracted out so as to reveal the seasonal part more clearly. In this and the other frames of Figure 3 a bar plot shows the relative magnitudes of the coefficients of the seasonal spectrum components $P_3(f)$ through $P_{13}(f)$ represented in the spectrum. One of these is negative.

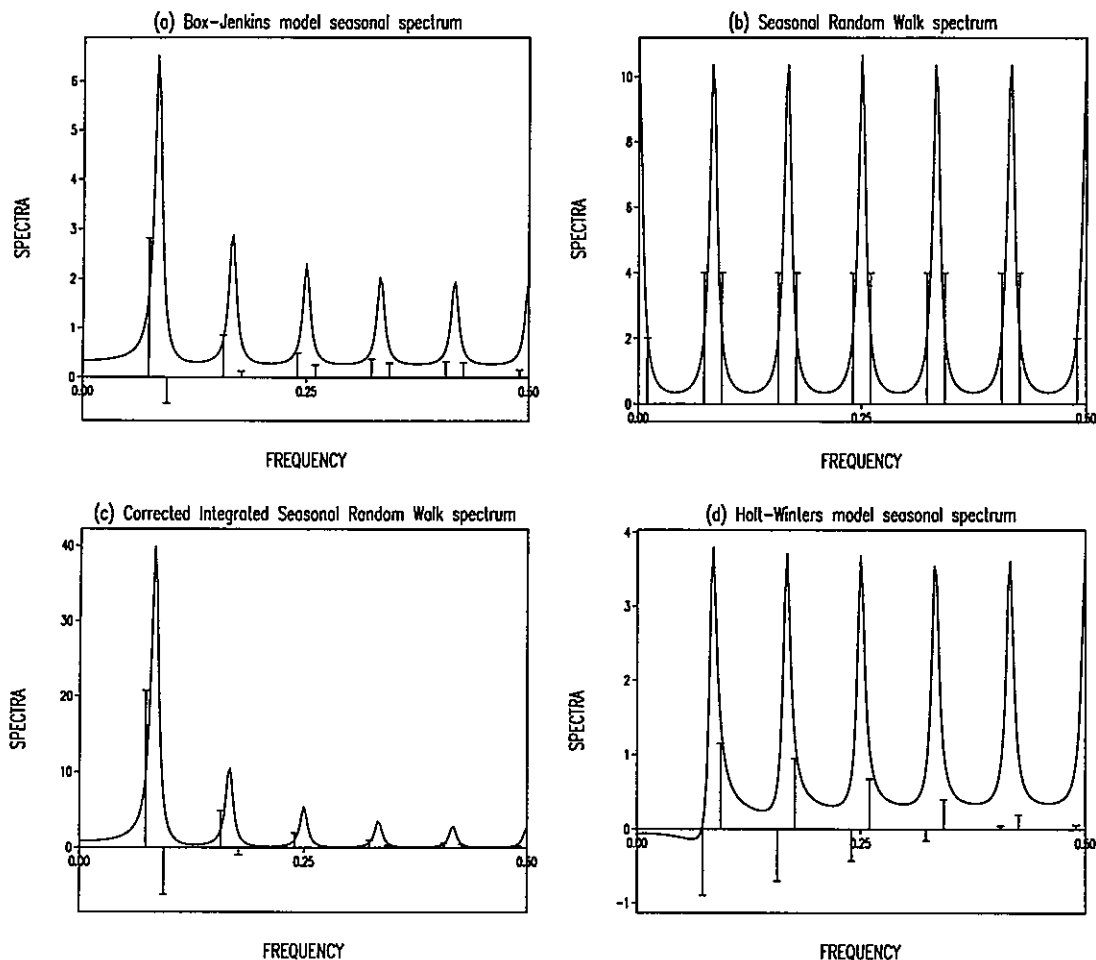


Figure 3: Model spectra, with bar plots indicating the coefficients of spectral components, for (a) A Box-Jenkins Model, (b) A Seasonal Random Walk, (c) A Corrected Integrated Seasonal Random Walk, (d) A Holt-Winters Model.

The spectrum of the 'airline' model can be expressed as the positive combination of four contributions. The first of these is the constant white noise spectrum, proportional to $P_0(f)$. The second is that of a random walk, proportional to $P_1(f)$, which has been removed in Figure 3a. The third contribution is the spectrum of a seasonal random walk $\nabla_{12}x_t = e_t$ which is proportional to $1/(1 - \cos 12\omega)$. This is shown in Figure 3b *without* the peak at frequency zero removed. All the peaks in this contribution are the same height and in fact all the ten seasonal components $P_3(f)$ through $P_{12}(f)$ have the same coefficient. The peaks at the ends of the range are proportional to

$P_1(f)$, the random walk component, and $P_{13}(f)$ each with *half* the coefficient of the other components. We shall later use this pattern, with the peak at zero removed, as a *combined* component to represent a purely seasonal feature of a model, referring to it as the *corrected seasonal random walk component*.

The fourth contribution to the airline model spectrum is the spectrum of an *integrated* seasonal random walk $\nabla\nabla_{12}x_t = e_t$. This models the falling away of the peaks towards higher frequencies and is shown in Figure 3c, but with the peak at frequency zero removed. We shall use this as another combined component to represent seasonality, referring to it as the *corrected integrated seasonal random walk component*. It is a combination of $P_3(f)$ to $P_{13}(f)$ in which the coefficients of $P_4(f)$ and $P_6(f)$ are negative and that of $P_8(f)$ is zero. After adding the contribution from the seasonal random walk just one seasonal coefficient remains negative in Figure 3a.

Figure 3d shows the spectrum corresponding to the Holt-Winters predictor with discount factors tuned for the logarithm of the airline series, given, in the notation used by Newbold (1988), by $A = 0.35$, $B = 0.01$ and $C = 0.75$. Again the peak at frequency zero due to $P_1(f)$ and $P_2(f)$ is removed. The pattern is very different to that seen in Figure 3a. The one-step ahead forecast error variance for the Holt-Winters predictor is only about 5% higher than that using the Box-Jenkins model, but the implications for evolving seasonality are seen to be dissimilar.

One other combined seasonal component is that in the basic structural model of Harvey and Todd (1983), which also appears in Newbold's (1988) comparison. It is represented using the seasonal summation operator as $S(B)x_t = e_t$. Its spectrum is found from that of the seasonal random walk in Figure 3b by multiplying by the gain $2(1 - \cos\omega)$ of the differencing operator ∇ , in contrast to *dividing* by this gain which furnishes the spectrum of the integrated seasonal random walk in Figure 3c. Consequently its spectrum has a pattern of increasing peaks towards high frequencies which we have not seen in the series we have modelled, and we do not illustrate it. Later structural models as described by Harvey (1990) allow components which separately model each seasonal peak. These components are equal to the *sum* of the pairs of seasonal components we describe together with the single seasonal component. A further model used by Harvey assumes a common coefficient for all these peaks. The result is the same as the corrected seasonal random walk component shown in Figure 3b, apart from a scaling factor and the relative doubling of the single seasonal component.

We stated earlier that if the coefficients of the components in our model are positive, then they can be interpreted as arising from unobserved independent components of the observed series. These series components can then be estimated; in particular they can be used to extract the seasonal components. The state space formulation we present in section 5 can be used for this purpose. However, as the foregoing discussion of Figure 3 indicates, it is quite possible to have some of the coefficients of individual components negative yet for their sum to represent a well defined positive component. Part of the value of fitting our component model is to be able to investigate whether all the coefficients are positive, and if not whether a particular combination of the fitted components, such as the set of seasonal components, is itself positive. As Burman (1980) shows, the Box-Jenkins 'airline' model always has a positive decomposition into non-seasonal and seasonal components. Figure 3d shows that this is not true for the Holt-Winters predictor.

The pairs of seasonal components which we have chosen are helpful in such an investigation. The first (second) component of each pair is zero at the upper (lower) frequency limit so that their combined contribution is positive if and only if both coefficients are positive. The coefficient of one component may be negative provided the sum of their coefficients is positive, which ensures that a positive peak results from their combination. A negative peak necessarily implies an inadmissible spectrum model; the estimation method may, rarely, imply this but the components would not be statistically significant and would be removed from the model. Similar constraints apply to other components, in that the coefficient V_2 of the integrated random walk and V_{13} of the single seasonal component must be positive.

Our general model is useful for checking the extent to which the imposition of ‘desirable’ constraints, such as replacing each seasonal component pair by its sum, or forcing all the seasonal components to have the same coefficient, affects the fit of the model. It also readily allows the fitting of subsets of components as in ordinary regression if the aim is to identify the sources of variability in level, trend and seasonality or simply to obtain a parsimonious model.

4 Model estimation using a modified sample spectrum

We now propose a method for fitting the linear model (8), selecting those terms in the model which are supported by the data.

For a stationary time series a spectrum model $S(f)$ may be fitted to the sample spectrum $S^*(f)$ of the observed series using what is generally known as the Whittle likelihood. This approximates the likelihood using the assumption that at the frequencies $f_\ell = \ell/n$, where n is the length of the observed series, $S^*(f)$ are independent values of Exponential variables with means $S(f)$. When the spectrum model is linear as in our case this is a standard generalised linear model, see McCullagh and Nelder (1983). The GLIM software may be used to fit the model, as in Diggle (1990), but there is always a possibility that this will fail when a linear model is fitted with the identity link due to negative fitted values arising in the iteratively reweighted least squares (IRLS) cycle of the method. This is one of the difficulties which our method overcomes.

Harvey (1990) has used the Whittle likelihood for fitting the variance coefficients in seasonal time series models. Because the likelihood is defined only for a stationary model he uses the sample spectrum of the differenced series w_t defined in (5) and the ‘differenced’ components. For our model these are obtained by multiplying all the components in (8) by the differencing gain in the denominator of (7) which reduces them all to polynomials in $C = \cos(2\pi f)$.

Using the sample spectrum of w_t has several drawbacks. Because the differencing gain has zeros at all the peak frequencies the visual effect is that the peaks in the sample spectrum of the data are invariably replaced by troughs at the same points in the spectrum of the differenced series. The evidence for a particular spectrum peak is then the, generally small, amount by which this differenced spectrum exceeds zero in the corresponding trough. If the component associated with a particular spectrum peak has a zero coefficient and is not in fact present in the model then the differenced model spectrum will have an exact zero at that point. Testing for such a zero component has inherent difficulties. It is a non-standard inference problem for the exponential error model, equivalent to the unit-root testing problem in the time domain. Moreover there is poor relative accuracy (which is the important measure of accuracy in exponential error models), close to this zero, of the differenced sample spectrum as an estimate of the model spectrum. The sample spectrum *must* be positive even if the model spectrum is zero. This can lead to the inclusion of a component which is not evident as a peak in the undifferenced spectrum. If a particular component is deemed to be absent then the remaining components should be refitted to the spectrum of a series obtained by applying a modified differencing operator. This operator is obtained by removing the factor associated with the deleted component from $\nabla\nabla_{12}$. The model components must be correspondingly modified.

The implication is to seek a method of fitting the model directly to the undifferenced series spectrum. For example the dynamic harmonic regression models of Ng and Young (1990) have since been fitted to the logarithms of an autoregressive spectral estimate of the undifferenced data, emphasising the importance of the flanks of the spectral peaks for determining the variance coefficients.

We now describe the method we recommend for overcoming the difficulties which arise from the (seasonal) unit root non-stationarity of (6); ie the factors $[1 - \exp(2\pi i F_j)B]$ of $\nabla\nabla_{12}$ which generate

the infinite spectral peaks at frequencies F_j . We argue that, provided the deterministic trend and seasonality terms of model (1) are included in the model, these unit roots may be modified to make the model stationary by replacing the above factors by $[1 - s \exp(2\pi i F_j)B]$. The value of the *shrinkage* s is chosen to be a little less than unity. Such near unit roots were proposed in Hannan's (1964) model of stochastic seasonality. The use of exact unit roots has since been favoured, see Hannan (1967) for example, partly because it is difficult to discriminate statistically between a root of unity and one slightly smaller, and also because the unit root model conveniently and robustly models the persistence of typical seasonal patterns without the explicit introduction of deterministic terms for trend and seasonality.

However, once *having* corrected for deterministic terms, the peaks we see, for example in Figure 1b, are quite consistent with the peaks of the *shrunk* stationary model.

The first plank of our estimation method is therefore to modify all the models we use by *shrinking* the unit roots (or rather, the reciprocals of these roots) in these models by the factor s . The implementation of this is described in the next section, but for example this is achieved for the seasonal IMA model (5) by multiplying the coefficient of B^j by s^j . We discuss the choice of s later. By this means the components $P_k(f)$ in (8) are modified to have finite peaks corresponding to stationary processes. Formulae for these *shrunk* components are given in the next section. All the illustrations of spectrum components in Figures 2 and 3 are of *shrunk* components with $s = 0.97$.

These *shrunk* components could then be fitted directly to the sample spectrum of the error series n_t of model (1), formed after correction for the deterministic trend and seasonality, so that the difficulties described above, of using the spectrum of the differenced series w_t , are avoided. We expect the estimated coefficients of the components to be insensitive to the precise choice of s just less than unity.

In practice we introduce another modification, so the second plank of our estimation method is to use the sample spectrum $W^*(f)$ defined in (4) of the residuals e_t from (3) rather than the errors n_t from model (1) for fitting our component model. The key point here is that the generalised differencing partially *whitens* the spectrum. This both reveals better the features that are to be modelled and gives better data to which the model components are fitted, by reducing leakage from the dominant low frequencies. Of course this requires that the model components $P_k(f)$, besides being *shrunk*, be also multiplied by the gain of the generalised differencing, or whitening, operator $(1 - \rho B)$ which is

$$w(f) = (1 + \rho^2 - 2\rho C). \quad (9)$$

We shall refer to this as fitting on the whitened scale and will call the new *shrunk* and whitened components $Q_j(f)$.

The third and final plank in our method is to avoid problems with the IRLS scheme of fitting used for the exponential error model. We do this, following de Jong (1985), by obtaining a consistent estimate of the final weights used in this scheme. We can then use just one step of (weighted) OLS regressions to obtain our estimates and carry out tests to select which components are supported by the data. These weights are in fact estimates of the fitted values of the model corresponding to the whitened residual spectrum $W^*(f)$. We obtain them by fitting the general model (5) using standard time domain exact likelihood methods which are very quick and reliable. The weight function $W(f)$ is then given by the spectrum defined in (6), but first modified by coefficient shrinkage in both numerator and denominator and then multiplied by the generalised differencing gain:

$$W(f) = \left| \frac{\theta(sB)}{(1 - sB)(1 - s^{12}B^{12})} \right|^2 \sigma_e^2 w(f). \quad (10)$$

If any of the components of (8) are not present, this will correspond to unit root factors in the

moving average operator of (6) which cancel unit root factors in the differencing operator of that equation. Such cancellations give rise to no problems in theory, and none that we have found in practice when estimating (5), provided exact likelihood estimation is used. Applying coefficient shrinkage to both differencing and moving average parts when calculating the weight function $W(f)$ ensures that any such cancellation is preserved and no spurious peaks are introduced.

To select the shrinkage parameter s and validate the weight function as a fit to the whitened sample spectrum we consider a discrete set of possible values $s = .99, .98, \dots$, calculate the corresponding value of $W(f)$ and plot it superimposed on the whitened sample spectrum $W^*(f)$. The visual comparison helps to confirm the adequacy of the fit and indicates if too much shrinkage (too low a value of s) has been used. A formal measure of the fit is the likelihood, or equivalently the deviance (minus twice the log likelihood relative to the saturated model), calculated as

$$\text{dev} = 2 \sum \left[\frac{W^*(f_\ell)}{W(f_\ell)} - \log \frac{W^*(f_\ell)}{W(f_\ell)} - 1 \right]. \quad (11)$$

The sum is taken over the frequencies $f_\ell = \ell/n$ for integer ℓ , *excluding* frequency zero and the seasonal frequencies F_j because regression upon the fixed trend and seasonality will have reduced the fitted values at these points to near zero. In practice we have found that the deviance usually improves (decreases) as s initially moves away from 1, which supports the use of some shrinkage. The estimation of the model coefficients however proves not to be very sensitive to the precise choice of s and we suggest that it not be reduced below 0.95.

The fitted values obtained using $s = 0.98$ for the inventory series are shown superimposed on the whitened spectrum in Figure 1b.

We now use the consistent estimator $W(f)$ as the weights in OLS regression for fitting the (shrunk) component model by taking as the 'response' and 'regressors':

$$\frac{W^*(f)}{W(f)}, \quad \frac{Q_j(f)}{W(f)}; j = 0, \dots, 13 \quad (12)$$

The frequencies used are again $f_\ell = \ell/n$ excluding frequencies zero and F_j . Note that the whitening factor $w(f)$ appears in both numerator and denominator of the regressors so may be cancelled for calculating this ratio.

We start by fitting all the 14 components. The maximum likelihood estimates of the coefficients in this case are a direct transformation of the 13 moving average parameters and residual variance in the fitted seasonal IMA model (5). This direct transformation can be obtained simply by making the response in the above OLS regression identically equal to 1. The coefficients obtained by carrying out a regression with the response as given should be very close to these.

Having fitted the full model we then apply the usual procedures of OLS regression to select a submodel with significant coefficients supported by their standard error estimates and the visual evidence of fit to the sample spectrum. The standard errors given by the weighted OLS are asymptotically valid and are useful for indicating important terms in the model. However, we recommend the use of deviance reduction tests, as presented by McCullagh and Nelder (1983), for decisions on their inclusion. These tests are similar to sum of squares reduction tests but correspond to likelihood ratio tests for fitting the model spectrum. The spectrum fit must however be positive for the deviance to be defined.

We shall also compare such subset component models with models in which fixed combinations of the components are used. Examples are the corrected seasonal random walk and the corrected integrated seasonal random walk components defined in section 3.

5 The model components and their representation

Provided that all the coefficients are positive in the component model (8) it is convenient to use a state space model to define the components; as for example in Harvey (1990). Each coefficient then appears naturally as the variance of an independent white noise input to this model. We shall however also give specific formulae for the model spectrum components for direct use in the estimation procedure just described. These are best expressed in terms of the variable $C = \cos(2\pi f)$, defined earlier, and constants $C_j = \cos(2\pi F_j)$, $S_j = \sin(2\pi F_j)$.

The state space representation has the form:

$$y_t = hx_t + \epsilon_{0,t} \quad (13)$$

where the observation vector h is a row of 13 constants combining the elements of the state vector x_t . The observation error $\epsilon_{0,t}$ gives rise to the white noise component $V_0 P_0(f)$ of the model, so that $P_0(f) = 1$ for all frequencies.

The states evolve according to the structural transition equation

$$x_t = Tx_{t-1} + \epsilon_t \quad (14)$$

where each element of ϵ_t is an independent white noise input, the k -th component having variance $\text{Var}(\epsilon_{k,t}) = V_k$. The transition matrix T which determines the structure of the model is block diagonal, the states being considered in pairs except for the final *single* state. In order that the states defined by this model follow the stationary processes used in our estimation procedure, rather than unit root processes, the transition matrix T should be multiplied by the shrinkage factor s . However we omit this for simplicity in the following presentation.

The contribution of the first pair of states is defined by the sub-model

$$\begin{bmatrix} h_1 & h_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \end{bmatrix} \quad ; \quad \begin{bmatrix} x_{1,t} \\ x_{2,t} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix}. \quad (15)$$

Then $\epsilon_{1,t}$ gives rise to a simple random walk and $\epsilon_{2,t}$ to an integrated random walk. Note that the state $x_{1,t}$ is the sum of these two and is the only state which contributes to the observation.

When the shrinkage factor is introduced the random walks become respectively an AR(1) with coefficient s and a repeated root AR(2) process, the formulae for the corresponding model spectrum components being

$$P_1(f) = \frac{1}{1 + s^2 - 2sC} \quad ; \quad P_2(f) = \left(\frac{s}{1 + s^2 - 2sC} \right)^2. \quad (16)$$

Note that the factor s^2 arises in the numerator of $P_2(f)$ as a consequence of the way the state transition is formulated; it may be omitted without any consequence of substance in the modelling.

The next five pairs of states are associated with the five seasonal periods F_1 through F_5 . The contribution of each pair is defined for $j = 1, \dots, 5$ by the sub-model

$$\begin{bmatrix} h_{2j+1} & h_{2j+2} \end{bmatrix} = \begin{bmatrix} \alpha_j & \beta_j \end{bmatrix} \quad ; \quad \begin{bmatrix} x_{2j+1,t} \\ x_{2j+2,t} \end{bmatrix} = \begin{bmatrix} C_j & -S_j \\ S_j & C_j \end{bmatrix} \begin{bmatrix} x_{2j+1,t-1} \\ x_{2j+2,t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_{2j+1,t} \\ \epsilon_{2j+2,t} \end{bmatrix} \quad (17)$$

where $\alpha_j = \sin(\pi F_j)$ and $\beta_j = \cos(\pi F_j)$. The formulae for the corresponding (shrunk) model spectrum components arising from $\epsilon_{2j+1,t}$ and $\epsilon_{2j+2,t}$ are

$$P_{2j+1}(f) = \frac{(1 - C_j)(u + sC)}{4[(uC_j - sC)^2 + (vS_j)^2]} \quad ; \quad P_{2j+2}(f) = \frac{(1 + C_j)(u - sC)}{4[(uC_j - sC)^2 + (vS_j)^2]}. \quad (18)$$

To simplify this expression we have defined $u = (1 + s^2)/2$ and $v = (1 - s^2)/2$. These are the components shown in Figure 2 using $s = 0.97$. When no shrinking is applied $s = 1$, $u = 1$, $v = 0$ and the components simplify further to

$$P_{2j+1}(f) = \frac{(1 - C_j)(1 + C)}{4(C_j - C)^2} \quad ; \quad P_{2j+2}(f) = \frac{(1 + C_j)(1 - C)}{4(C_j - C)^2}. \quad (19)$$

Both components defined in (18) have an ARMA representation of the form

$$(1 - 2C_j sB + s^2 B^2)y_{j,t} = z_{j,t} = (1 - \theta sB)e_{j,t} \quad (20)$$

where for the first component $\theta = 1$ and $\text{Var}(e_{j,t}) = \alpha_j^2$. For the second, $\theta = -1$ and $\text{Var}(e_{j,t}) = \beta_j^2$.

We comment upon the choice of the observation vector elements α_j and β_j , dropping the subscript j for simplicity. Let us allow any choice of the values of these elements, any choice of the variances V_α, V_β of $\epsilon_{2j+1,t}, \epsilon_{2j+2,t}$ which appear in this model *and* any choice of the covariance $V_{\alpha\beta}$ between these two inputs. Then the total contribution to the model of the state equations (17) is a series component which also has an ARMA representation of the form (20) and whose spectrum is of a similar form to those in (18). The denominator is exactly the same; the numerator is of the general form $\gamma_0 + 2\gamma_1 C$ which is the spectrum of $z_{j,t}$, the moving average part of this model. The five quantities $\alpha, \beta, V_\alpha, V_\beta$ and $V_{\alpha\beta}$ determine then only two free coefficients γ_0 and γ_1 of the spectrum contribution. It is therefore usual to constrain the five quantities in some way, and a commonly used approach is to set $\alpha = 1$ and $\beta = 0$. In fact the pair of states can always be redefined (rotated and scaled) to ensure that this holds, without *any* change to the state equations. There is a corresponding rotation and scaling of the inputs but this requires merely a re-interpretation of V_α, V_β and $V_{\alpha\beta}$. We give here two reasons to justify our different choice of α and β ; further details are given by Haywood (1994).

Firstly, it is widely agreed that it is desirable to avoid using the covariance $V_{\alpha\beta}$; doing so leaves just the required number of two free coefficients, the variances V_α and V_β which will have their associated spectrum components contributing to the total model spectrum. However, for the frequency $F_3 = 1/4$ this does *not* lead to two free coefficients γ_0 and γ_1 when using $\alpha = 1, \beta = 0$. In fact the two spectrum components associated with V_α and V_β are then proportional, in the ratio $1 : s^2$. Using the choice of α and β which we propose this is not the case, as is seen directly from (18). Our choice leads to precisely two linearly independent components for each of these seasonal frequencies. If in fact $V_{\alpha\beta}$ were taken as non-zero using our choice of α and β it would lead to a further component of the same form as in (18) but with numerator simply $2vS_j$. This, conveniently, vanishes when shrinkage is not applied, in which case $V_{\alpha\beta}$ does not enter the equation.

Secondly, when shrinkage is not applied, *any*, necessarily non-negative, spectrum component generated by a state sub-model such as (17) can be expressed as a combination of the two components in (18) using *positive* values of V_α and V_β . With other choices of observation vector elements such as $\alpha = 1$ and $\beta = 0$ this is not always possible. This point is not important if it is merely required to fit a spectrum model, but is important if it is desired to implement a state-space model which corresponds to that fitted model and use it for forecasting or seasonal adjustment. The explanation of the point lies in the fact that the components we define lie at the extremes of possible positive spectra, the first (second) component being zero at the upper (lower) limit of the frequency range. This does not hold precisely when shrinkage is applied but the difference is very small, proportional to $(1 - s)^2$. This point is made in the context of the particular state sub-model (17). We have previously pointed out that using the components we define in (18) to fit a general spectrum pattern, negative coefficients can result which prevents the use of the Kalman Filter for estimation.

A final point to note is that models in which it is assumed that $V_\alpha = V_\beta$ are fitted using the sum of the components in (18). This sum is in fact the same as that obtained by using $\alpha = 1$ and

$\beta = 0$. When the inputs are of equal variance and uncorrelated, this property is unchanged by rotating them.

One further sub-model requires definition, that for the last, single, seasonal component which is

$$\begin{bmatrix} h_{13} \end{bmatrix} = \begin{bmatrix} 1 \end{bmatrix} \quad ; \quad \begin{bmatrix} x_{13,t} \end{bmatrix} = \begin{bmatrix} -1 \end{bmatrix} \begin{bmatrix} x_{13,t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_{13,t} \end{bmatrix}. \quad (21)$$

Including the shrinkage factor, this is an AR(1) model with coefficient $-s$. It give rise to a spectrum component

$$P_{13}(f) = \frac{1}{1 + s^2 + 2sC} \quad (22)$$

which mirrors at the upper frequency limit of 0.5 the AR(1) (shrunken random walk) component at the lower limit of 0.

6 Modelling examples

Our first example series is illustrated in Figure 1a. Figure 1b shows the sample spectrum on the whitened scale together with the fit obtained from the seasonal IMA model. Using the procedure described in section 4 we fit a series of component models to the sample spectrum. This procedure requires only ordinary least squares, which provides approximate standard errors for estimates of the component coefficients. As explained in section 4, we use the (residual) deviance defined in (11) and deviance differences, rather than the residual sum of squares, for carrying out model comparisons. The residual deviance for the (shrunken) seasonal IMA model is 140.11. On the assumption that the fitted model is correct, this statistic should come from a distribution with mean $1.154p - 14 = 152$ and variance $1.29p - 14 = 172$ where $p = 144$ is the number of spectrum points used to fit the model. A Normal approximation is adequate for this number of points and in this case there is no evidence of lack of fit of the seasonal IMA model. This is not a very powerful test of model adequacy but is certainly a necessary qualification for proceeding further, given that our method is based on the assumption that this fit provides a consistent estimate of the spectrum.

We first fit the full 14 component model (model 1) described in section 3, with the variation described there, that each seasonal pair of components is replaced by its sum and difference. This reduces collinearity. Table 1 shows the estimated coefficients and their standard errors for this model. The fit of this model appears identical to that shown in Figure 1b. It is important to refit the model to obtain the baseline deviance for comparison with, and testing of submodels. It is in fact slightly reduced to 135.8; the change is attributable to the use of the shrunken components and the switch from time domain to frequency domain estimation methods. The model is equivalent to the shrunken seasonal IMA model. From the t values in Table 1 there is no evidence to support inclusion of the seasonal pair difference terms. Neither is there evidence to support the white noise or the integrated random walk component. However we retain these last two in the model until we have determined the best form of the seasonal components.

For model 2 we therefore simply remove the seasonal pair difference components. Because the coefficients of the terms retained are little changed, these are not shown. The fit of this model is shown in Figure 4a. The deviance increased by only 0.7, to 136.5. Since five parameters are removed, the appropriate test statistic distribution is chi-squared with five degrees of freedom. Hence exclusion of all seasonal pair difference terms is supported strongly, and the model reduces to a nine parameter structural model with different variances for each harmonic seasonal pair sum.

We now consider models with a single component for the seasonality. For model 3 we use the corrected seasonal random walk component for this purpose and for model 4 the corrected integrated seasonal random walk. These were defined in section 3.

white noise 21.0 (0.11)	random walk 1709.0 (4.45)	1st seasonal pair sum 48.4 (1.81)	2nd seasonal pair sum 23.19 (2.94)	3rd seasonal pair sum 10.03 (2.05)	4th seasonal pair sum 3.60 (2.32)	5th seasonal pair sum 1.65 (0.94)
integrated random walk 2.62 (1.05)	single seasonal component 0.444 (0.59)	1st seasonal pair diff 101. (0.49)	2nd seasonal pair diff 47. (0.46)	3rd seasonal pair diff -19.4 (-0.30)	4th seasonal pair diff -20.7 (-0.77)	5th seasonal pair diff -4.9 (-0.36)

Table 1: Spectrum component coefficient estimates for models of the series of unfilled orders for newspapers with (t) values.

Comparing model 3 with model 1 the deviance increase of 28.2 for a saving of 10 degrees of freedom lies at the upper 0.2% point; hence there is strong evidence against model 3.

If we note how, in Table 1, the coefficients of the seasonal components fall away from the lower to the higher harmonics, we obtain an indication that model 4 would better represent the seasonality by a single component and this is the case. The deviance increase compared to model 1 is only 4.9 on 10 degrees of freedom, (well below the 50% point).

The t values for the white noise and integrated random walk components were respectively -0.25 and 0.93. Removing these two terms had no adverse effect so that the final model included just two components, the random walk and the corrected integrated seasonal random walk. Their respective coefficients (t values) were 1864.1 (9.70) and 584.4 (3.70), while the residual deviance was 142.09. The fit of this model on the whitened scale is shown in Figure 4b. Whatever comparisons are made with the other models there is no strong statistical evidence to reject this model, which we recommend as the best for this series.




Figure 4: Spectrum models for the monthly series of unfilled orders.

Our second example is another US macroeconomic series – wholesale sales of sporting goods from January 1967 to February 1989 inclusive. This is plotted in Figure 5a, after taking natural logarithms and multiplying by 1000. The sample spectrum of the series on the whitened scale is

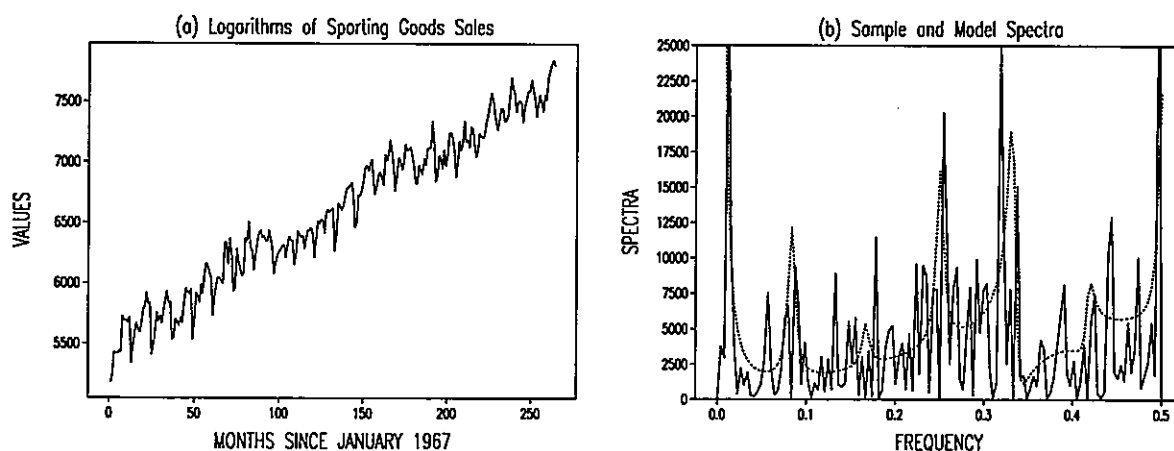


Figure 5: (a) Monthly values of the logarithms of sporting goods sales, (b) Sample and model spectra for the series.

shown in Figure 5b with the fit of the seasonal IMA model superimposed. The coefficient ρ used in the generalised differencing was 0.655 for this series, the shrinkage factor s used for the components was 0.97. The residual deviance for the seasonal IMA model fit was 144.35 using $p = 124$ frequency points giving a tolerable Normal deviate of 1.24 using a mean of 129.1 and a variance of 146.0. A simple correction for trading day effects has been made to the series by removing first a cycle at frequency 0.3477 which was evident as a sharp peak in the original whitened spectrum. This frequency is well known to be associated with trading day effects. One other, smaller, peak was also removed at the frequency $0.4310 = 1/12 + 0.3477$ which is associated with a modulation of the fundamental seasonal frequency $1/12$ by the trading day effect. Table 2 shows the estimated coefficients and their t values for model 1, the 14 component model. The fit was close to that shown in Figure 5b, but with some re-adjustment of coefficients giving a somewhat better residual deviance of 136.24. Because the coefficient of component 3, the integrated random walk, was negative, this component was removed for all subsequent analyses. The reason for this negative coefficient may be that the low frequency peak appears to lie a little above zero, so the random walk components may be attempting to approximate a low frequency business cycle in the data. Model 2 was determined by removing insignificant terms one by one, the terms remaining being indicated by a † in Table 2. The coefficients and standard deviations did not change appreciably; the fit for this model is shown in Figure 6a. The deviance for this model was 147.42 on 118 degrees of freedom, the increase of 11.2 from the 14 component model being reasonable for a saving of 8 degrees of freedom, corresponding to the upper 20% point of the appropriate chi-squared distribution.

Model 3 arose from attempts to simplify the seasonal components using both the corrected seasonal random walk component and the corrected integrated seasonal random walk component; only the latter was retained. The estimated coefficients for this model were 1868.7 ($t = 28.94$) for the white noise, 823.5 ($t = 17.88$) for the random walk, 702.3 ($t = 13.10$) for the corrected integrated seasonal random walk and 281.7 ($t = 21.95$) for the fourth seasonal pair difference. The fit of this 4 component model is shown in Figure 6b. The deviance increase from the model 1 was 17.9 on 10 degrees of freedom, which is at the upper 5.6% point. Figures 6a and 6b clearly reveal the nature of the choice between these two models - whether to go for a compromise in the more marginal model 3 or stay with the less parsimonious but better fitting model 2.

Note that for both models 2 and 3, one seasonal pair difference component was necessary in the fit. Without this the deviance increase is extreme. For model 2 the coefficient of the remaining term

white noise† 1708. (4.39)	random walk† 668. (2.62)	1st seasonal pair sum† 67.0 (1.97)	2nd seasonal pair sum 7.30 (0.89)	3rd seasonal pair sum† 22.1 (2.05)	4th seasonal pair sum† 16.99 (3.17)	5th seasonal pair sum 3.68 (0.14)
integrated random walk -0.87 (-0.77)	single seasonal component 6.07 (1.14)	1st seasonal pair diff 12. (0.66)	2nd seasonal pair diff 34. (0.24)	3rd seasonal pair diff 104. (0.64)	4th seasonal pair diff† 377.5 (3.90)	5th seasonal pair diff -17.1 (-0.44)

Table 2: Spectrum component coefficient estimates for models of the series of sporting goods sales with (t) values.

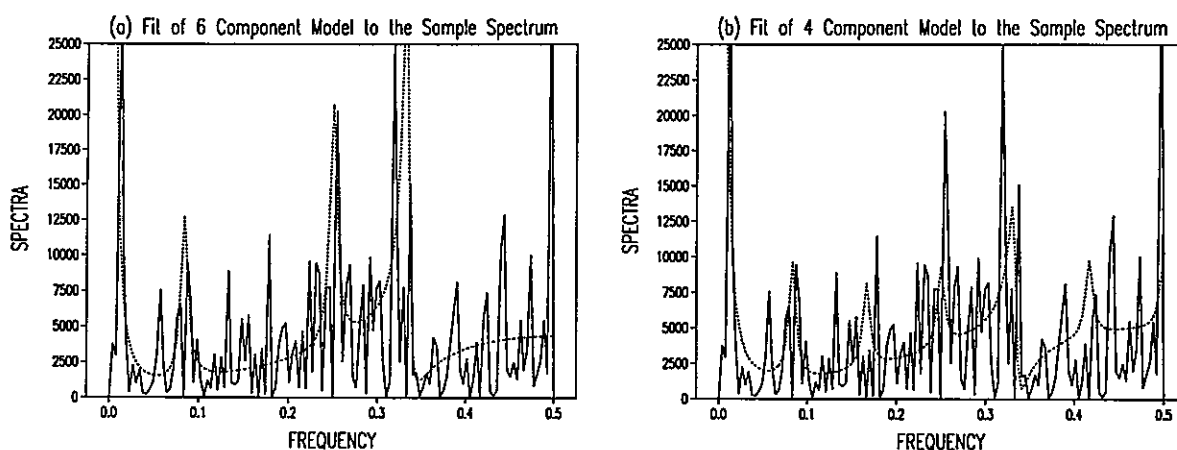


Figure 6: Further spectrum models for the series of sporting goods sales.

associated with that seasonal peak falls from 17.46 ($t=3.37$) to -1.07 ($t=-0.32$), so that this peak, which appears clearly in Figure 6a, is not fitted at all. For model 3, the use of a single component with peaks at all seasonal frequencies protects against this type of lack of fit. Its coefficient falls only from 702.3 ($t=13.10$) to 592 ($t=5.12$).

We have seen other examples similar to this one, in which a seasonal peak appears quite skewed, with an appreciable fall in the spectrum between that seasonal frequency and the next. Use of a seasonal pair difference term then appears necessary to obtain a good fit to the spectrum and in particular to avoid a poor estimate of the seasonal pair sum whose coefficient measures the amplitude of the peak. The disadvantage is that the seasonal pair difference component is necessarily negative over some frequency range so can only be interpreted in conjunction with other terms. Only if the coefficient of this component is smaller (in absolute value) than that of the corresponding seasonal pair sum component is their total contribution positive and representable by (17). This is not the case in this example. It does however become positive when the white noise component is also added. For this example that would leave the random walk component as the only non-seasonal feature of the model. We emphasise therefore that limiting the models to those with only positive coefficients in order to permit the use of the Kalman Filter can lead, for some time series, to a poor fit of the whole model. Conversely, the method we propose gives a

viable alternative to estimation by the Kalman Filter even if a structural model is required.

7 Discussion

We have presented a procedure which allows a flexible model for evolving trend and seasonality to be readily investigated for a given time series. This model may be applied to forecasting and seasonal adjustment, the latter being the application for which careful assesment of evolving seasonality is perhaps most important, and the interpretation of the components most valuable for constructing adjusted series. The first example shows how a parsimonious model may be readily identified using this method. The second example provides justification for considering all the terms in the model lest important seasonal features are poorly estimated.

The model building process we describe may be related to procedures for seasonal unit root testing, for example those of Canova and Hansen (1995) and more recently those of Hylleberg and Pagan (1997) who have returned to the evolving seasonality models presented by Hannan, Terrell and Tuckwell (1970). In our model components we distinguish however between the condition of a unit root, ie $s = 1$ as opposed to $s < 1$, and the presence or absence of a stochastic peak, ie $V_k > 0$ or $V_k = 0$. Our view is that establishing whether s is equal to or just less than unity is relatively unimportant and that inference for V_k is both more important and relatively insensitive to the choice of s . We choose a value for s just less than unity in order to carry out that inference, for V_k , in a robust and revealing way in the frequency domain. We favour this approach rather than that of *fixing* $s = 1$ and *then* supposing that testing for V_k is the same as testing for the presence of a seasonal unit root.

By assuming that $s < 1$ we ensure that we are carrying out inference upon a stationary model thus avoiding some of the non-standard asymptotic inference problems for unit root models. For the finite sample case, however, the distinction between standard and non-standard inference is illuminated by our frequency domain methods. When we apply our least squares method to test for the presence of a shrunken spectrum component, the important statistic is the combination of spectral ordinates in the neighbourhood of the seasonal frequency, weighted by the values of the seasonal peak component at that frequency. Asymptotically the central limit theorem will apply to this weighted combination of exponential (χ_2^2) variates, but for typical seasonal series which are modelled its distribution will be appreciably skewed. The effect on the distribution of the test statistic could be computed and would be similar to that found for the non-standard distributions associated with unit root tests. This point is made to acknowledge that no slight of hand associated with unit root shrinkage can magic away the non-standard testing problem. The finite sample properties of the deviance reduction tests that we have used could be checked by simulation studies.

Finally we discuss procedures for applying the results of the modelling process to forecasting and seasonal adjustment. If the fitted model can be represented in terms of the components $P_j(f)$ all with positive coefficients, or in terms of combinations of these such as the integrated seasonal random walk which has a simple state space representation, then state space methods can be applied directly for forecasting and seasonal adjustment. If some coefficients are negative, as we have shown may be the case, then a final form SARIMA representation of the model can be constructed using available tools, and Wiener filtering used to separate the trend and seasonal parts. We leave open the question of whether to retain the shrinkage factor at the application stage. Our limited experience suggests a slightly better model fit is achieved using modest shrinkage, which suggests that the fixed trend and seasonal terms be retained in the model. Once the model is estimated however, it is possible to check whether removing the shrinkage still gives a positive spectrum fit and then to apply the seasonal unit root forms of the model, state space or ARIMA, without the need for the deterministic trend and seasonal terms.

Acknowledgements

The first author was supported by a research studentship of the UK Science and Engineering Research Council during the early stages of this work. The second author gratefully acknowledges the financial assistance and hospitality of the School of Economics and Finance, Victoria University, Wellington in support of his visit there to complete this work.

References

- Box, G.E.P. and Jenkins, G.M. (1976). *Time Series Analysis, Forecasting and Control*, revised edn. Oakland: Holden-Day.
- Bruce, A.G. and Jurke, S.R. (1996). Non-Gaussian seasonal adjustment: X-12-ARIMA versus robust structural models. *Journal of Forecasting*, 15, 305–328.
- Burman, J. P. (1980). Seasonal adjustment by signal extraction. *Journal of the Royal Statistical Society A*, 143, 321–337.
- Canova, F. and Hansen, B.E. (1995). Are seasonal patterns constant over time? A test for seasonal stability. *Journal of Business and Economic Statistics*, 13, 237–252.
- Diggle, P.J. (1990). *Time Series: A Biostatistical Introduction*. Oxford: Oxford University Press.
- de Jong, P. (1985). An asymptotically efficient estimator for variance component time series models. *Methods of Operations Research*, 50, 275–286.
- Hannan, E.J. (1964). The estimation of a changing seasonal pattern. *Journal of the American Statistical Association*, 59, 1063–1077.
- Hannan, E.J. (1967). Measurement of a wandering signal amid noise. *Journal of Applied Probability*, 4, 90–102.
- Hannan, E.J., Terrell, R.D. and Tuckwell, N.E. (1970). The seasonal adjustment of economic time series. *International Economic Review*, 11, 24–52.
- Harrison, P.J. and Stevens, C.F. (1976). Bayesian forecasting. *Journal of the Royal Statistical Society B*, 38, 205–247.
- Harvey, A.C. (1990). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- Harvey, A.C., and Todd, P.H.J. (1983). Forecasting economic time series with structural and Box-Jenkins models [with discussion]. *Journal of Business and Economic Statistics*, 1, 299–315.
- Haywood, J. (1994). A Frequency Domain Investigation of Model Based Prediction. Unpublished PhD thesis. Lancaster: Lancaster University.
- Hylleberg, S. and Pagan, A. R. (1997). Seasonal integration and the evolving seasonals model. *International Journal of Forecasting*, 13, 329–340.
- McCullagh, P., and Nelder, J.A. (1983). *Generalised Linear Models*. London: Chapman and Hall.
- McKenzie, E. (1976). A comparison of some standard seasonal forecasting systems. *The Statistician*, 25, 3–14.
- Newbold, P. (1988). Predictors projecting linear trend plus seasonal dummies. *The Statistician*, 37, 111–127.
- Ng, C.N. and Young, P.C. (1990). Recursive estimation and forecasting of non-stationary time series. *Journal of Forecasting*, 9, 173–204.

WORKING PAPER SERIES

The main purpose of this series is to reach a wide audience quickly for feedback on recently completed or in progress research. All papers are reviewed before publication.

A full catalogue with abstracts and details of other publications is available, for enquires and to be included in our distribution list, write to:

Angela Dolan
Manager Administration Services
GSBGM, Victoria University of Wellington,
PO Box 600, Wellington, New Zealand
Tel: (04) 471 5366; Fax: (04) 496 5454

Code in bold denotes order number, eg: WP 1/91

--- Group denotes the author's academic discipline Group (note this does not necessarily define the subject matter, as staff's interests may not be confined to the subjects they teach).

1990 - 1994 titles available on request.

- WP 1/95** **Management Group**
Gilbertson, D.K., Wright, H., Yska, G, Gilbertson, D.W. and 1994 Students of MGMT 306 'Kiwi entrepreneurs: A study.'
- WP 2/95** **Management Group**
Cavana, R. 'Policy issues related to coastal and international shipping in New Zealand'
Shipping policy issues. *Transportant: The Journal of the Chartered Institute of Transport in New Zealand* 1995, Vol 25, No 2, 17-19.
- WP 3/95** **Information Systems Group**
Bonner, Marcus 'On seeing information systems as bridges'
- WP 4/95** **Management Group**
Cavana, Bob, Rob Crozier, Barrie Davis and Perumal Pillai 'A survey of academic staff attitudes towards the system of academic titles used in New Zealand universities'
Attitudes Towards Reclassifying Academic Titles in New Zealand Universities. *The Journal of Higher Education Policy and Management* 1996, Vol 18, No 1.
- WP 5/95** **Econometrics Group**
Krawczyk, J.B. and G. Zaccour 'Pollution management through levies and subsidies'
- WP 6/95** **Marketing Group**
Ashill, Nicholas and Malcolm Wright 'Marketing information systems - A review and reconceptualisation'
- WP 7/95** **Information Systems Group**
Casey, Mary-Ellen 'An exploratory study into the use of information technology as an important enabler of organisational differentiation in the financial sector'
- WP 8/95** **Economics Group**
Boles de Boer, David and Lewis Evans 'The economic efficiency of telecommunications in a deregulated market: the case of New Zealand'
- WP 9/95** **Management Group**
Mabin, Victoria J. 'Using spreadsheet optimisation facilities as a decision aid within the theory of constraints framework'
- WP 10/95** **Economics Group**
M. Khaled, M.D. Adams and M. Pickford 'Estimates of scale and scope economies in the New Zealand life insurance industry.'
- WP 11/95** **Economics Group**
John A. Carlson and Robert A. Buckle 'Price duration with two-sided pricing rules'
In Karl Heinrich Oppenlaender and Guenter Poser (Eds.) *Business Cycle Surveys: Forecasting Issues and Methodological Aspects*. Aldershot, Avebury. 1996: 101-118.

- WP 12/95** **Economics Group**
Ganesh Nana 'Developing a multi-sectoral CGE model of the New Zealand economy.'
- WP 13/95** **Money and Finance Group and Economics Group**
Stephen Burnell, Lewis Evans and Shuntian Yao 'Network games: The optimal network contract and the efficiency of bypass in oligopolistic network industries under light regulation'
- WP 14/95** **Economic History Group**
Gordon Boyce 'The Nickel Syndicate, 1901 - 1939'
- WP 15/95** **Money and Finance Group**
Jan Whitwell 'Monetary disinflation with inflation inertia: Central bank autonomy in an open economy'
- WP 16/95** **Economics Group**
Emery, H. Daniel V. Gordon and Doug McClintock 'On the efficacy of construction site safety inspections.'
- WP 17/95** **Economics Group**
Ganesh Nana 'An inter-temporal CGE model with fiscal and sector balances'
- WP 18/95** **Economics Group**
Jie Zhang 'Government debt, human capital, and endogenous growth'
- WP 19/95** **Accountancy Group**
Zahirul Hoque and Manzurul Alam 'Quality management and accounting in a New Zealand service organisation: Towards an institutional perspective on management accounting'
- WP 20/95** **Economics Group**
Paul Calcott 'Can we test for supplier-induced demand by comparing informed with uninformed consumers?'
- WP 1/96** **Management Group**
M.W. Lee, J. Bennett, R.J. Taylor and R.Y. Cavana 'A dynamic simulation model for possum and gorse control on a farm woodlot.'
R.Y.Cavana, M.W.Lee, J.Bennett & R.J.Taylor. 1996. Possum and gorse control on a farm woodlot in New Zealand: A system dynamics analysis.
Asia-Pacific Journal of Operational Research, Vol 13, No. 2, pp. 181-207.
- WP 2/96** **Economics Group**
Jie Zhang 'Optimal public investments in education, and endogenous growth'
- WP 3/96** **Economics Group**
Paul Tompkinson 'The benefits of tariff reductions in the presence of psychological adjustments costs.'
- WP 4/96** **Economics Group**
Shuntian Yao 'A note on the decision of a sales maximizer in response to the increase of per unit cost.'
- WP5/96** **Economics Group**
Karen Palmer and Margaret Walls 'Optimal policies for solid waste disposal: taxes, subsidies and standards'
- WP6/96** **Economics Group**
Margaret Walls and Jean Hanson 'Distributional impacts of an environmental tax shift: the case of motor vehicle emission taxes.'
- WP7/96** **Economics Group**
Lloyd S Shapley and Shuntian Yao 'Dynamic properties of the Nash equilibrium.'
- WP 8/96** **Accountancy Group**
Kamran Ahmed, Kazi Feroz Alam and Manzurul Alam 'An empirical study of factors affecting accounting students' career choice in New Zealand'
- WP9/96** **Accountancy Group**
Andrew MC Smith and Paul V Dunmore 'Tax Avoidance and the Financial Structures of Non -Resident Controlled Companies in New Zealand'
- WP10/96** **Econometrics Group**
L Fraser Jackson. Relative Prices and Inflation.'
- WP11/96** **Public Policy Group**
Robert Gregory. 'Reserve Bank independence, political responsibility, and the goals of anti-democratic policy: A political 'Cri de Coeur' in response to an economist's perspective.'
- WP12/96** **Economics Group**
Robert A. Buckle and John A. Carlson. 'Inflation and asymmetric price adjustment.'

WP13/96

Econometrics Group

J.B. Krawczyk, O. Pourtallier and M. Tidball. 'Modelling and solution to the municipal effluent management problem.'

WP14/96

Economics Group

Viv B. Hall and David Rae. 'Fiscal expansion, interest rate risk premia, and wage reactions: Some illustrative simulations with NBNZ-DEMONZ.'

WP15/96

Economics Group

Viv. B. Hall, Kunhong Kim and Robert A. Buckle. 'Pacific rim business cycle analysis: Dating, volatility and synchronisation.'

WP16/96

Management Group

S. Davenport, C. Grimes and J. Davies. 'Collaboration and organisational learning: A study of a New Zealand collaborative research programme.'

WP1/97

Information Systems Group

Marcus Bonner and Mary-Ellen Casey. 'The informed individual - a consideration of importance to the business organisation.'

WP2/97

Management Group

R.Y. Cavana, I.G. Harrison, F.E.B. Heffernan and C.C. Kissling. 'Freight transport industry in New Zealand.'

WP3/97

School of Economics and Finance

Steffan Berridge and Jacek B Krawczyk. 'Relaxation Algorithms in finding Nash Equilibria.'

WP 4/97

School of Economics and Finance

Margaret Walls and Karen Palmer. 'Upstream Pollution, Downstream Waste Disposal, and the Design of Comprehensive Environmental Policies'

WP 1/98

School of Economics and Finance

John Haywood and Granville Tunnicliffe Wilson. 'Selection and Estimation of Trigonometric Component Models for Seasonal Time Series'