

A STRUCTURAL MODEL WITH INTERVENTIONS FOR NEW ZEALAND SAWN TIMBER PRODUCTION

DONGWEN LUO, GEOFFREY JONES, AND JUDITH DENNIS

Received 16 January 2005; Revised 16 May 2005; Accepted 17 May 2005

Timber production in New Zealand was privatized in 1987. We examine the effects of this change on the level of New Zealand sawn timber production, and changes in the seasonal pattern, using a state-space model with intervention variables. We describe the formulation and estimation of the state-space model, and show how it can be used to examine both the structural changes around the time of privatization and the gradually evolving seasonal pattern in production. We also show how the model can be used to forecast future production.

Copyright © 2006 Dongwen Luo et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. Introduction

The Statistics Section of the Ministry of Agriculture and Forestry, New Zealand, collects statistics on sawn timber production both annually (effectively a census) and quarterly (a sample survey). Quarterly data is published provisionally, and final figures are given when the annual survey has been completed. Figure 1.1 shows the quarterly figures for the period from March 1973 to June 2002 (118 values). The series shows a strong seasonal pattern with wave crest at quarter 3 (September) and trough at quarter 1 (March). The amplitude of the seasonal variation does not seem to increase as the level of the series increases, which suggests additive rather than multiplicative seasonality. The apparent dip in the series in the late 1980s coincides with the disestablishment of the New Zealand Forest Service and the sale or corporatisation of state forest assets (Brown and Valentine [2]; Clarke [4]; Rhodes and Novis [17]). Thereafter there appears to be a dramatic increase in trend. In addition, it is thought that the seasonal pattern has changed over time, although this is not thought to be related to privatization. There are several factors affecting the seasonality: weather, which reduces logging operations in a wet season; regular holidays; and annual maintenance of sawmilling operations. Change in the seasonal pattern is expected because of technological improvements, in particular the extended use of kiln drying to allow for an increase in production through the autumn and winter.

2 New Zealand sawn timber production

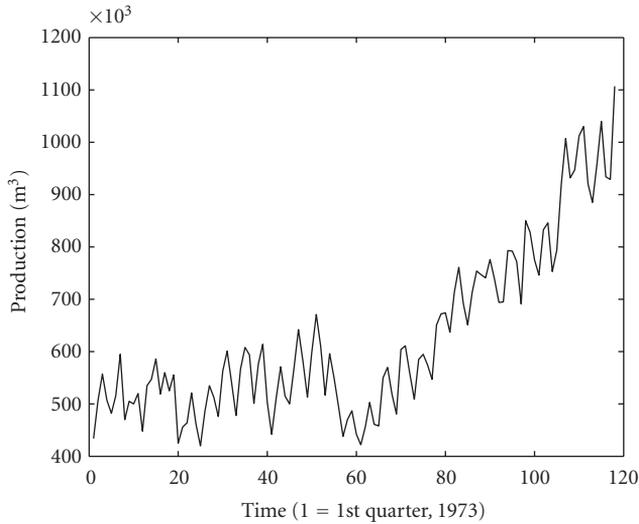


Figure 1.1. Total NZ sawn timber production, quarterly from March 1973 to June 2002. Privatization occurred at the end of 1986 (time period 57).

The aims of our analysis are to investigate the effects of privatization and to examine this changing seasonality. This requires an appropriate decomposition of the observed series into trend, seasonal, and other components, allowing for the structural change during the privatization episode. Our approach is to use the framework of the basic structural model, a form of state-space model, with intervention variables to account for the sudden dip in the level of the series and the simultaneous increase in trend.

We first give a brief review of the basic structural model in Section 2, then demonstrate in Section 3 the use of interventions in the basic structural model to incorporate the effect of privatization. Section 4 considers alternatives to estimation of the variance parameters. We investigate in Section 5 the ability of the model to forecast future values by comparing the forecasts for the next two years with the available data and with the forecasts from an ARIMA model. Section 6 summarizes and discusses our findings.

2. Basic structural model

The basic structural model (Harvey [8]) decomposes an observed time series Y_t into trend, seasonal, and random components as follows:

$$Y_t = \mu_t + S_t + e_t, \quad (2.1)$$

where μ_t represents the underlying level of the series, S_t the seasonal factor, and e_t a random error term. The level μ_t is modelled as a “random walk with drift”:

$$\mu_t = \mu_{t-1} + \beta_{t-1} + u_{1t}, \quad (2.2)$$

where the slope β is allowed to evolve over time according to a random walk:

$$\beta_t = \beta_{t-1} + u_{2t}. \quad (2.3)$$

The seasonal factors can also evolve over time by means of a “seasonal random walk”:

$$S_t = -S_{t-1} - S_{t-2} - S_{t-3} + u_{3t}. \quad (2.4)$$

The four error series e_t , u_{1t} , u_{2t} , u_{3t} represent random shocks to the series, and are assumed here to be independent and normally distributed with variances σ_e^2 , σ_{u1}^2 , σ_{u2}^2 , σ_{u3}^2 , respectively. The variances determine the smoothness of each component: if a variance is small, then that component changes slowly over time. For example, if σ_{u2}^2 is close to zero, then the local slope β_t will be almost constant over time, implying a more or less fixed trend in timber production; a larger value for this variance would give more flexibility and allow the slope to vary, so that future forecasts would be based more on the recent trend in production. Values for these variance parameters can be estimated from the data using the method of maximum likelihood, or chosen to give appropriately smooth components.

The model described above can be written in matrix form as a “state-space model,” in which the “observation equation”

$$Y_t = \mathbf{A}\mathbf{x}_t + e_t \quad (2.5)$$

tells us how the observed series is related to the underlying state of the system, and the “state equation”

$$\mathbf{x}_t = \mathbf{\Phi}\mathbf{x}_{t-1} + \mathbf{u}_t \quad (2.6)$$

shows how the state evolves over time. The state here represents the level μ , slope β , and seasonal components S . In the basic structural model as given above, the state vector would be $\mathbf{x}_t = (\mu_t, \beta_t, S_t, S_{t-1}, S_{t-2})^T$, the observation matrix $\mathbf{A} = (1, 0, 1, 0, 0)$, the state innovation $\mathbf{u}_t = (u_{1t}, u_{2t}, u_{3t}, 0, 0)^T$, and the state transition matrix

$$\mathbf{\Phi} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & -1 & -1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}. \quad (2.7)$$

The likelihood of the hyperparameters σ_e^2 , σ_{u1}^2 , σ_{u2}^2 , σ_{u3}^2 can be constructed by running the Kalman filter (Kalman [13] Schweppe [18]). Maximum likelihood estimates can then be obtained by using a quasi-Newton optimizer, or alternatively the EM method of Shumway and Stoffer [20]. The components of the state vector \mathbf{x}_t are then estimated by running the Kalman smoother (de Jong [5]). This overall procedure can be regarded as a model-based alternative to more pragmatic methods, such as X-11 (Shiskin et al. [19]) and the HP filter (Hodrick and Prescott [12]) for decomposing an observed series into trend and other components.

4 New Zealand sawn timber production

Structural models with a cyclical component have been considered by Watson [22] and Harvey and Jaeger [11]. A stochastic cyclical component C_t , with a variable period, can be modelled using the Box-Jenkins AR(2) model

$$C_t = \phi_1 C_{t-1} + \phi_2 C_{t-2} + u_{4t}, \quad (2.8)$$

so we expand the state vector to $\mathbf{x}_t = (\mu_t, \beta_t, C_t, C_{t-1}, S_t, S_{t-1}, S_{t-2})^T$ and the observation matrix to $\mathbf{A} = (1, 0, 1, 0, 1, 0, 0)$. The state equation is now

$$\begin{pmatrix} \mu_t \\ \beta_t \\ C_t \\ C_{t-1} \\ S_t \\ S_{t-1} \\ S_{t-2} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_1 & \phi_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \mu_{t-1} \\ \beta_{t-1} \\ C_{t-1} \\ C_{t-2} \\ S_{t-1} \\ S_{t-2} \\ S_{t-3} \end{pmatrix} + \begin{pmatrix} u_{1t} \\ u_{2t} \\ u_{3t} \\ 0 \\ u_{4t} \\ 0 \\ 0 \end{pmatrix}. \quad (2.9)$$

In our analysis of NZ timber production, we begin with the basic structural model without a cyclic component, and later extend our analysis to investigate the possible presence of cycles.

3. Intervention analysis

Intervention analysis in time series refers to the modelling of a sudden change in the nature of the series, and can be of two forms. In the first, a known event at a particular point in time is investigated for its effect, if any, on the series. For example, Harvey and Durbin [10] examined the effect of seat belt legislation on deaths by car accident in the UK. In the second form, the existence and location of possible shocks is not known in advance, and the problem is akin to outlier detection (de Jong and Penzer [6]). In practice we may require something between these two extremes. For the NZ timber data we know the date of privatization (Q1-1987, or $t = 57$), but structural changes may be expected to occur either just before or just after this date.

Interventions can be introduced into the basic structural model via the observation equation (2.5), the state equation (2.6), or both. In our case, inspection of the raw series suggests a drop in the level (μ) component and an increase in the slope (β) component, so we replace (2.6) with

$$\mathbf{x}_t = \mathbf{\Phi} \mathbf{x}_{t-1} + \lambda \otimes \mathbf{I}_t + \mathbf{u}_t, \quad (3.1)$$

where “ \otimes ” denotes componentwise product, $\lambda = (\lambda_\mu, \lambda_\beta, 0, 0, 0)^T$, and \mathbf{I}_t is a 5×1 vector of zeros except at the intervention time(s), when either the first or second component, or both, is one. \mathbf{I}_t here represents the pattern of the interventions (when they occur) and λ the sizes of the shocks introduced to the system. Luo [14] showed how the method of

Table 3.1. Effect of interventions on the likelihood for the basic structural model.

Interventions	$-2\log L$	AIC
None	983.59	991.59
Level only	974.89	984.89
Slope only	979.54	989.54
Level and slope	971.40	983.40

Shumway and Stoffer [20] using the EM algorithm can be extended to simultaneously estimate both λ and the variance parameters, given \mathbf{I}_t .

To determine the form of \mathbf{I}_t , that is, where to place the level and slope interventions, we considered two complementary strategies. The method of de Jong and Penzer [6] is based on examining the smoothed values of the state interventions $E(\mathbf{u}_t | Y_1, \dots, Y_n)$ which are obtained from the Kalman smoother fitted to a null model without interventions. Here we find an unusually large negative value for u_{1t} (level component) at $t = 55$, two quarters before the nominal date of privatization, and a number of large positive values for u_{2t} (slope component) beginning at $t = 59$. An alternative strategy is to fit the full model including interventions, estimating all parameters including the locations of the interventions. This can be very computationally intensive, but can be achieved using Markov chain Monte Carlo methods (Carter and Kohn [3]). Our approach here is to search a grid of values in the neighbourhood of $t = 57$ and compare the values of the log-likelihood, maximized over the variance parameters. This confirms the best placing of the interventions as $t = 55$ for the level and $t = 59$ for the slope.

Table 3.1 gives the value of $-2\log$ -likelihood, minimized over the variance parameters, and the Akaike information criterion AIC (Akaike [1]), for models with and without interventions at the above locations. Here AIC has been calculated as $-2\log L + 2k$ where k is the number of estimated parameters. The model with both level and slope interventions has the lowest AIC so is the preferred model.

Figure 3.1 shows selected components from the final run of the Kalman smoother with level and slope interventions. The estimates from the full model are $\sigma_e^2 = 0$, $\sigma_{u1}^2 = 637$, $\sigma_{u2}^2 = 0$, $\sigma_{u3}^2 = 58.0$, $\lambda_\mu = -87.5$, and $\lambda_\beta = 9.31$. Because the slope error variance is zero, the slope component is piecewise constant and so the level μ is modelled as a random walk with drift, the drift increasing from 1.03 per quarter before privatization to 10.34 per quarter afterwards. This corresponds to an annual average increase of about 4.1 thousand cubic metres before, and 41.4 after, privatization. Thus the model suggests that a sudden, and considerable, drop in production of about 87 thousand cubic metres was experienced leading up to privatization, but that subsequently production has been increasing at a steady rate.

Figure 3.1(b) shows the changing seasonal pattern. The seasonal difference between Quarter 2 (June) and Quarter 3 (September) seems to have diminished, whereas Quarters 4 and 1 (December and March) have initially diverged and then reconverged. Seasonal adjustment can be accomplished easily in the state-space framework, by subtracting the estimated seasonal factors from the observed series. In our model the seasonally adjusted

6 New Zealand sawn timber production

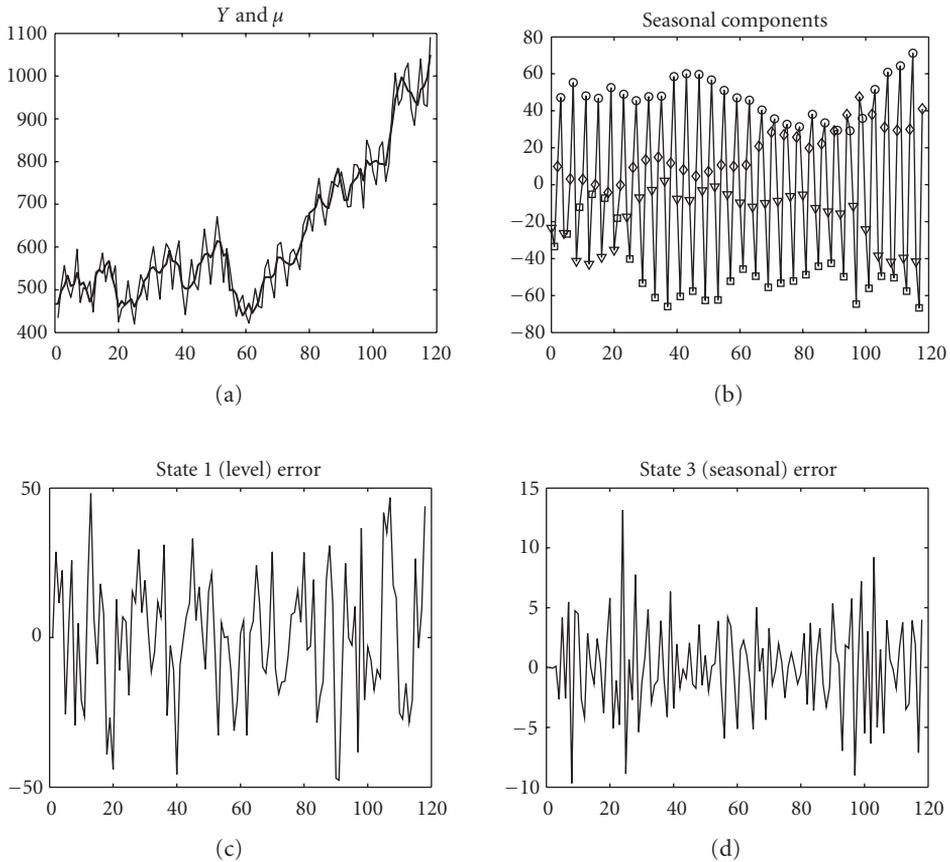


Figure 3.1. Fitted structural model showing raw and smoothed series (a); seasonal factors (b); level (μ) component residuals (c); seasonal component residuals (d). The seasonal factor legend is Q1 = square, Q2 = diamond, Q3 = circle, Q4 = triangle.

series consists simply of the level component, which is plotted in Figure 3.1(a) along with the actual series.

Figures 3.1(c), 3.1(d) show the estimated error series for the level and seasonal components, and can be used for model checking. The seasonal error series suggests a few periods of rapid change in the seasonal pattern, in 1974–1975, 1978–1979, and 1997–1998, perhaps corresponding with technological changes. The level error series appears to be positively autocorrelated, and inspection of the level component suggests a possible medium-term cycle, perhaps a business cycle, of about 4 to 5 years. Allowing for this could give a smoother level component and a different partitioning of the total variance in the series, since in the current model the cycles could be inflating the level variance.

The maximum likelihood estimates obtained from the structural model with cyclical component are $\phi_1 = 1.027$, $\phi_2 = -0.286$, $\sigma_e^2 = 22.6$, $\sigma_{u1}^2 = 0$, $\sigma_{u2}^2 = 434.3$, $\sigma_{u3}^2 = 0$, $\sigma_{u4}^2 = 73.9$, $\lambda_\mu = -109.4$, and $\lambda_\beta = 8.24$. Figure 3.2 shows selected components from the final

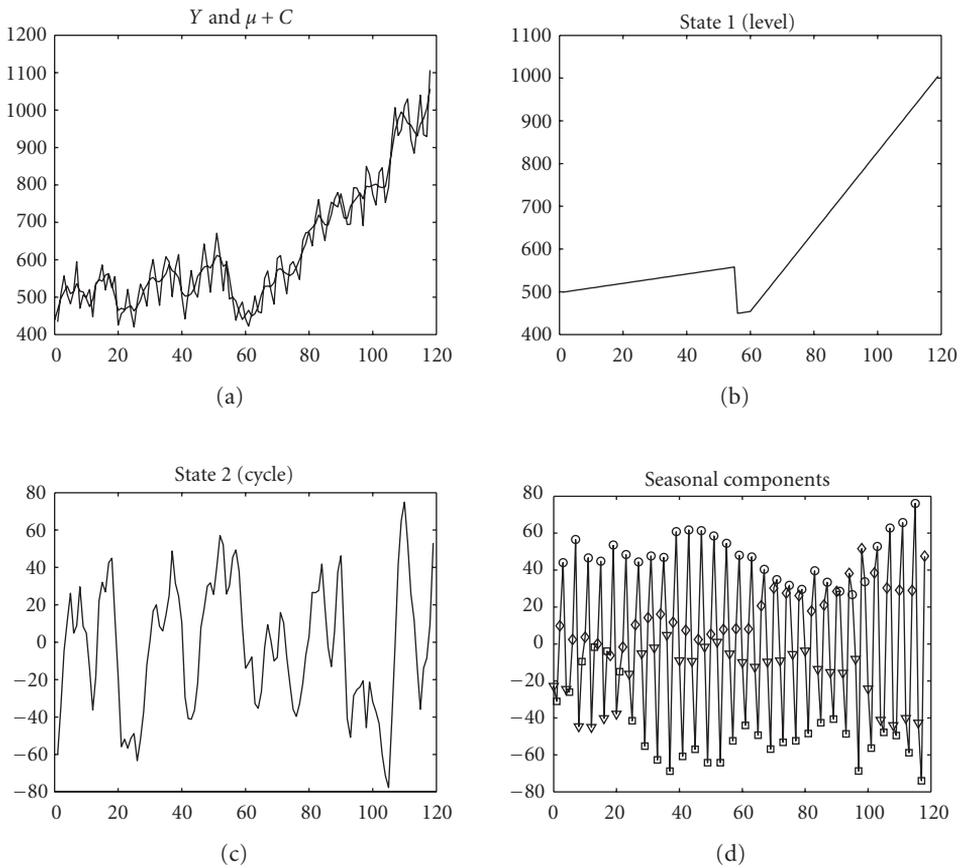


Figure 3.2. Fitted model with cyclical component showing raw and smoothed series (a); level (μ) component (b); cycle (C) component (c); seasonal component (d). The seasonal factor legend is Q1 = square, Q2 = diamond, Q3 = circle, Q4 = triangle.

run of the Kalman smoother. Again the slope error variance is zero, but now the error variance of the level μ also drops to zero, so the estimated level is composed of a series of line segments as shown in Figure 3.2(b). The cyclical component has absorbed most of the variation, so that instead of a smooth cycle it is a rather noisy, but stationary, AR(2). The estimated values of ϕ_1 , ϕ_2 imply a spectrum which decreases monotonically with increasing frequency, so that there is no dominant cycle. The minimized value of $-2\log$ -likelihood is now 966.31, which is only a slight improvement on the previous model, and the AIC is slightly higher at 984.31. The pattern of changing seasonality as shown in Figure 3.2(d) is very similar to that of the previous model. Thus there seems to be little to be gained from adding the cyclical component. The only practical consequence of the addition would be in forecasting future values. The simpler model projects the level forward as a straight line and adds the most recent seasonal factor to give the forecast. The additional cyclical component, by contrast, will incorporate recent departures from

the straight line trend via the AR(2) model. This might be considered to be appropriate on heuristic grounds.

4. Imposing smoothness

As an alternative to maximum likelihood estimation of the hyperparameters, that is, the variances, we now consider a more heuristic approach, namely, the choice of values which produce what is considered to be an appropriate amount of smoothness in each component. This is equivalent to the use of preset smoothing parameters in exponential smoothing or Holt-Winters decomposition, rather than choosing the values which optimize some fitting criterion.

In modelling New Zealand timber production, we might decide for forecasting purposes to have a smoothly changing slope, so that it is the “local slope” which is projected forwards in the state equation (2.6). We might also choose, a priori, a smooth cyclical component to represent “the business cycle.” We can impose the required smoothness by setting the corresponding variances to suitably small values. An informal trial-and-error approach, with visual inspection of the smoothness of the estimated component series, can be used to obtain reasonable parameter values. Figure 4.1 shows the result of fitting the model with $\sigma_{u1}^2 = 20$, $\sigma_{u2}^2 = 20$, $\sigma_{u3}^2 = 0.1$. The estimates of the other parameters are $\phi_1 = 1.712$, $\phi_2 = -0.862$, $\sigma_e^2 = 489.5$, $\sigma_{u4}^2 = 25.4$, $\lambda_\mu = -102.3$, and $\lambda_\beta = 7.46$. This is in some respects a very different model, with most of the variance being incorporated into the observation error. The changes in the seasonality are in broad agreement with the earlier models, but are now much smoother. The AR(2) model for the cyclical component now has a dominant frequency, that is, a peak in the spectrum, at $\cos^{-1}(-\phi_1(1 - \phi_2)/4\phi_2) = 0.391$ (Harvey [9, page 181]), corresponding to a period of 16.1 quarters or about 4 years.

5. Forecasting

In forecasting future values from a state-space model, the state equation (2.6) allows us to project the state forward in time, setting future values of \mathbf{u}_{t+h} equal to zero. The observation equation (2.5) then converts these into forecasts \hat{Y}_{t+h} . This can be done by a simple modification of the Kalman filter, which can also produce prediction errors for the forecasts; Harvey [8] or Shumway and Stoffer [21]. For illustration we use the model from Section 3, that is, without a cyclical component. Since this model has a fixed trend component, the state equation projects the level forward as a straight line with constant slope and adds the most recent seasonal factor to give the forecast.

Figure 5.1(a) shows the forecasts from this model for the next 8 quarters, up to June 2004, and compares them with the most recent figures available from the Ministry of Agriculture and Forestry. The actual values can be seen to be reasonably close to the forecasts, and well within the prediction intervals, except for the first value, September 2002, for which the actual production figure is considerably higher than expected. Note that here we are forecasting all eight future values using only the data up to June 2002. In practice one could update future forecasts as new data become available, using the Kalman filter.

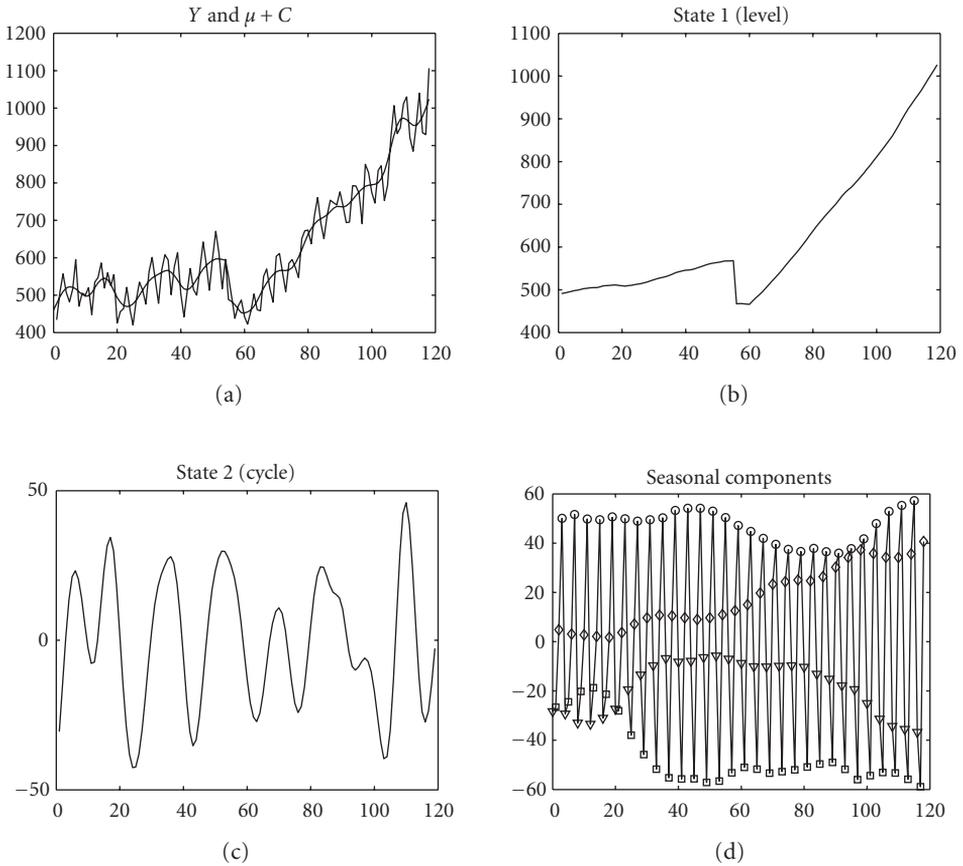


Figure 4.1. Fitted model with smooth cyclical component showing raw and smoothed series (a); level (μ) component (b); cycle (C) component (c); seasonal component (d). The seasonal factor legend is Q1 = square, Q2 = diamond, Q3 = circle, Q4 = triangle.

An $ARIMA(0, 1, 2) \times (2, 1, 0)_4$ model, selected using the AIC criterion, was estimated as

$$(1 + .45B^4 + .34B^8)(1 - B^4)(1 - B)Y_t = (1 - .09B + .15B^2)e_t. \quad (5.1)$$

The AIC for the fitted ARIMA model was 985.18, a little larger than that of the best state-space model. The forecasts from this model are shown in Figure 5.1(b). The predictions are more accurate in the short term but less accurate overall, with a mean absolute percentage error of 8.1% as compared to 5.1% for the state-space model. The prediction intervals are noticeably wider for the ARIMA forecasts. This is to be expected, as the innovations in the series resulting from the privatization episode will, if not modelled, inflate the forecast error variance.

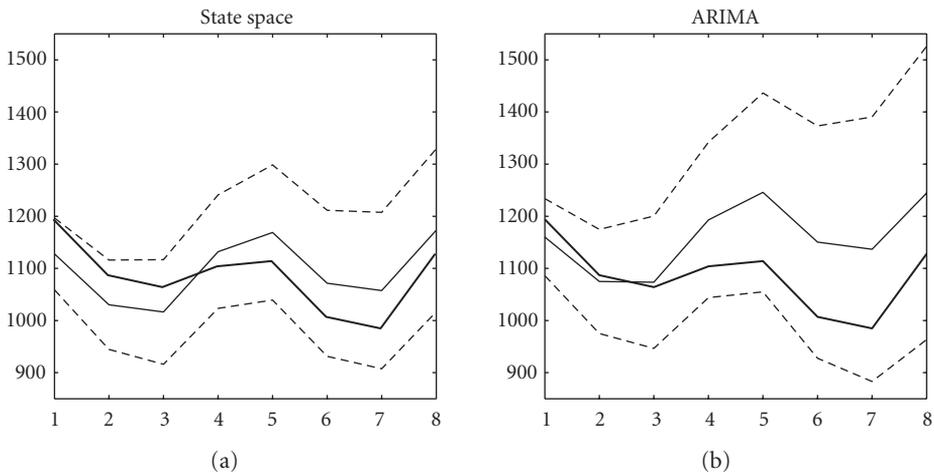


Figure 5.1. Forecasts for September 2002 to June 2004. The heavier solid line shows the actual values. The broken lines give the prediction intervals.

6. Discussion

We have used the framework of the basic structural model to decompose the New Zealand timber production series, including interventions in the structural components to allow for possible changes due to privatization. This allows us to estimate and visualize the way the seasonal pattern changes over time. We have chosen an additive model, based on the observation that the amplitude of seasonal variation appears to be approximately constant. It could perhaps be argued, a priori, that proportional seasonal variation would be more appropriate. If the above analyses are repeated with the log-transformed data, the seasonal components are found to converge gradually over time. This would imply that seasonal variability is decreasing, perhaps because of technological factors. The other implication of this alternative model is that production increased at a rate of approximately 0.8% before privatization, and 6% afterwards. It is interesting to note that the forecasts from this model are a little more accurate, but the prediction intervals considerably wider.

The inclusion of a cyclical component in the model is not really supported by the data. The existence of business cycles and the methods used to estimate them are contentious issues in econometrics. Quah [16] has pointed out an identifiability problem in the decomposition of a given series into trend and cyclical components without prior assumptions about smoothness. In the econometrics literature there has been some controversy over the “calibration,” as opposed to the estimation, of models (Hansen and Heckman [7]). Calibration here is taken to mean the use of values for some parameters based on historical reported values or on what is felt to be appropriate behaviour. If we believe in business cycles, we can impose restrictions on certain parameter values to produce cycles of the expected length of 4–6 years. But the apparent cycles in timber production could otherwise be explained using a random walk model, and there is no compelling statistical

evidence for preferring the former. The restrictions on certain parameters necessary to detect business cycles could perhaps be incorporated formally into the analysis via the Bayesian approach (Pole et al. [15]).

A basic ARIMA model actually performed a little better in short-range out-of-sample forecasting. The advantage of the state-space approach is its ability to separate, model, and describe the various components of a series.

Acknowledgment

We thank the editor and referee for their helpful and insightful comments.

References

- [1] H. Akaike, *A new look at the statistical model identification*, IEEE Transactions on Automatic Control **19** (1974), no. 6, 716–723.
- [2] C. L. Brown and J. Valentine, *The process and implications of privatization for forestry institutions: focus on New Zealand*, Unasylva **45** (1994), no. 178, 11–19.
- [3] C. K. Carter and R. Kohn, *On Gibbs sampling for state space models*, Biometrika **81** (1994), no. 3, 541–553.
- [4] M. Clarke, *Devolving forest ownership through privatization in New Zealand*, Unasylva **50** (1999), no. 199, 35–44.
- [5] P. de Jong, *Smoothing and interpolation with the state-space model*, Journal of the American Statistical Association **84** (1989), no. 408, 1085–1088.
- [6] P. de Jong and J. Penzer, *Diagnosing shocks in time series*, Journal of the American Statistical Association **93** (1998), no. 442, 796–806.
- [7] L. P. Hansen and J. J. Heckman, *The empirical foundations of calibration*, Journal of Economic Perspectives **10** (1996), no. 1, 87–104.
- [8] A. C. Harvey, *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge, 1989.
- [9] ———, *Time Series Models*, 2nd ed., Harvester Wheatsheaf, New York, 1993.
- [10] A. C. Harvey and J. Durbin, *The effects of seat belt legislation on British road casualties: a case study in structural time series modelling*, Journal of the Royal Statistical Society. Series A **149** (1986), 187–227.
- [11] A. C. Harvey and A. Jaeger, *Detrending, stylized facts and the business cycle*, Journal of Applied Econometrics **8** (1993), 231–247.
- [12] R. Hodrick and E. Prescott, *Post-war US Business cycles: an empirical investigation*, Journal of Money, Credit and Banking **29** (1997), no. 1, 1–16.
- [13] R. E. Kalman, *A new approach to linear filtering and prediction problems*, Transactions of the ASME - Journal of Basic Engineering **8** (1960), 95–108.
- [14] D. Luo, *Intervention analysis in the basic structural model*, unpublished Master of Applied Statistics thesis, Massey University, Palmerston North, New Zealand, 1999.
- [15] A. Pole, M. West, and J. Harrison, *Applied Bayesian Forecasting and Time Series Analysis*, Chapman & Hall, New York, 1994.
- [16] D. Quah, *The relative importance of permanent and transitory components: identification and some theoretical bounds*, Econometrica **60** (1992), 107–118.
- [17] D. Rhodes and J. Novis, *The Impact of Incentives on the Development of Plantation Forest Resources in New Zealand*, MAF Information Paper No: 45, 2002.
- [18] F. C. Scheweppe, *Evaluation of likelihood functions for Gaussian signals*, IEEE Transactions on Information Theory **11** (1965), no. 1, 61–70.

12 New Zealand sawn timber production

- [19] J. Shiskin, A. H. Young, and J. C. Musgrave, *The X-11 Variant of the Census Method II Seasonal Adjustment Program*, Technical Paper 15, US Department of Commerce, Bureau of the Census, Washington, DC, 1967.
- [20] R. H. Shumway and D. S. Stoffer, *An approach to time series smoothing and forecasting using the EM algorithm*, *Journal of Time Series Analysis* **3** (1982), 253–264.
- [21] _____, *Time Series Analysis and Its Applications*, Springer Texts in Statistics, Springer, New York, 2000.
- [22] M. W. Watson, *Univariate detrending methods with stochastic trends*, *Journal of Monetary Economics* **18** (1986), no. 1, 49–75.

Dongwen Luo: Fonterra Research Centre, 5301 Palmerston North, New Zealand
E-mail address: dongwen.luo@fonterra.com

Geoffrey Jones: Institute of Information Sciences and Technology, Massey University,
5301 Palmerston North, New Zealand
E-mail address: g.jones@massey.ac.nz

Judith Dennis: Innovation and Research, Ministry of Agriculture Forestry,
4000 Wellington, New Zealand
E-mail address: dennisj@maf.govt.nz