On growth and volatility regime switching models for New Zealand GDP data

Bob Buckle
New Zealand Treasury

David Haugh
New Zealand Treasury

Peter Thomson
Statistics Research Associates Ltd
New Zealand

March 7, 2002

Abstract

This paper reviews and documents methodology for fitting hidden Markov switching models to New Zealand GDP data. A primary objective is to better understand the utility of these methods for modelling growth and volatility regimes present in the New Zealand data and their interaction. Properties of the models are developed together with a description of the estimation methods, including use of the EM algorithm. The models are fitted to New Zealand GDP and production sector growth rates to analyse changes in the mean and volatility of historical business cycles. The paper discusses applications of the methodology to dating business cycles, identifies changes in growth performances, and examines the timing of growth and volatility regime switching between GDP and its production sectors. Directions for further development are also discussed.

JEL classification: C22 Time series models; E23 Production; E32 Business fluctuations, cycles; O47 Measurement of economic growth.

Keywords: Hidden Markov models; regime switching; growth; business cycles; volatility; production sectors; GDP.

1 Introduction

Interpretation of New Zealand's trend economic growth during the 1990s has been a central issue in recent debate concerning New Zealand's growth potential, its growth performance relative to that achieved in other developed economies and debate surrounding the impact of economic reforms. One of the difficulties is deciding on the interpretation that should be placed on a run of observed higher or lower growth rates. When should such a sequence be interpreted as a change in the mean growth rate or, for that matter, a change in volatility?

One of the purposes of this study is to obtain more timely and sensitive measures of changes in New Zealand's economic growth performance and to develop better methods for the identification of shifts in growth and volatility regimes. If successful, this will enhance interpretation of current data and policy analysis. These are important objectives given the data limitations that confront researchers measuring real economic growth in New Zealand and the relatively volatile nature of these data by comparison to those for large–scale developed economies such as the United States, Japan, and the larger European economies.

A common way to characterise the growth process and to interpret the stages of the business cycle is to assume that first differences of the logarithms (the growth rates) of real GDP follow a linear stationary process (using moving averages, or the filter suggested by Hodrick and Prescott (1980), or structural time series techniques such as those advocated by Harvey (1989)) and that optimal forecasts of GDP are assumed to be a linear function of their lagged values. However there is considerable evidence to suggest that departures from linearity are an important feature of many key macroeconomic series. This evidence includes the documentation of business cycle asymmetries by Neftci (1984) and Sichel (1987) and a growing body of research showing that real output responds asymmetrically to nominal demand shocks (Cover (1988); de Long and Summers (1988); Morgan (1993), Karras (1996)) and that inflation can induce an asymmetric real output response to changes in demand (See Rhee and Rich (1995) for US evidence; Olekalns, (1995) for Australian evidence; Buckle and Carlson (2000) for New Zealand evidence).

Such findings have prompted the development of time series models for GDP that assume that the growth rates follow a non–linear stationary process. An important development in this regard is the Hamilton (1989) model of the US business cycle. Hamilton assumes US GNP growth is subject to discrete shifts in regimes where the regimes are discrete episodes over which the dynamic behaviour of the series is markedly different. His approach is to use the Goldfeld and Quandt (1973) Markov switching regression to characterise changes in the parameters of an autoregressive process. The economy may be in a fast growth or slow growth phase with the switch between the two governed by the outcome of a Markov process. Regime switching models such as these have also been heavily used in many other disciplines including finance (Hamilton and Susmel (1994)), meteorology (Zucchini and Guttorp (1991)) and speech recognition (Rabiner (1989)) to name but a few.

Hamilton found that the best fit of his regime switching model to US GNP data gave growth regimes that were similar to NBER dating of business cycles. This suggests that this modelling approach could be used as an alternative objective algorithm for dating business cycles and, more generally, it opens up the possibility of capturing different dynamics during the different stages of the business cycle. Since Hamilton's model of the US business cycle, the Markov switching autoregressive model has become increasingly popular for the empirical characterisation of macroeconomic series. Several researchers have found this framework to be a useful approach for characterising business cycles including, for the US business cycle, Lam (1990), Boldin (1994), Durland and McCurdy (1994), Filardo (1994), Diebold and Rudebusch (1996), Kim (1994) while Krolzig (1997) has also found it a useful tool for investigating the business cycles of Australia, Canada, France, Germany, Japan and the United Kingdom.

The purpose of this paper is to develop and estimate Markov regime switching models for New Zealand real GDP growth and the growth of its component production sectors. The aim is to better understand how these models can be used to identify changes in growth and volatility in a small scale open economy with relatively short time spans of data. The success with which these types of models have been used to identify changes in growth and volatility in larger economies suggests they are worth exploring for New Zealand, notwithstanding the data difficulties.

Another principal reason for developing regime switching models is to explore the merits of a different way of thinking about how an economy's growth rate evolves and the interpretation to be placed on changes in the growth and volatility of real output. In effect, these models block the data into periods of time (regimes comprising a number of consecutive quarters) whose time evolution is directly modelled, in addition to the quarter–to–quarter evolution within regimes. Thus the various time scales in the data are separately modelled within a simple, open framework that should allow enhanced economic and policy analysis. Because of its readily understood structure, this type of analysis can also be used as an exploratory tool to help guide appropriate specification of other model based methods.

The remainder of the paper is structured as follows. Section 2 describes the hidden Markov switching model (HMM model) that we have fitted to New Zealand GDP growth data together with its specification and properties. Section 3 discusses issues concerning the estimation and fitting of HMM models. The results of fitting the HMM models to New Zealand real GDP data and to five production sectors are discussed in Section 4. Sections 5 and 6 apply these fitted models to the dating of turning points in the New Zealand business cycle and growth process, and to the comparison of the timing of cycles in production sectors and total GDP. Conclusions are drawn and directions for future research are discussed in Section 7.

2 Model

In general we assume that we have available observations Y_t ($t = 1, \ldots, T$) on some stationary macroeconomic time series where Y_t follows the general model

$$Y_t = \mu_{S_t} + \sigma_{S_t} X_t \qquad (t = 0, \pm 1, \ldots) \tag{1}$$

and the focus is on the case where Y_t represents the growth rates of GDP or one of its production sectors. The stochastic process S_t is an unobserved stationary finite Markov chain that takes on the values $1, \ldots, N$, which index the states of the system. Thus the level μ_{S_t} and the volatility σ_{S_t} switch between the N ordered pairs of values $(\mu_1, \sigma_1), \ldots, (\mu_N, \sigma_N)$ according to S_t. The stochastic process X_t is assumed to be a zero mean stationary Gaussian process which is independent of S_t and so

$$E(Y_t|S_t) = \mu_{S_t}, \qquad \text{Var}(Y_t|S_t) = \sigma^2_{S_t}\text{Var}(X_t)$$

give the time varying mean and variance of Y_t when the S_t are known. Moreover, given the S_t, the autocorrelation function of Y_t is the same as the autocorrelation function of X_t which is time invariant. In this context the unobserved components μ_{S_t} and σ_{S_t} represent a stochastic trend (location) and a stochastic volatility (scale) respectively.

Through the hidden states S_t the model allows for discretely changing levels and volatility over time. A simple example is given in Figure 1 where there are two states ($N = 2$) with $S_t = 1$ corresponding to a low level μ_1, low volatility σ_1 state, and $S_t = 2$ corresponding to a high level μ_2, high volatility σ_2 state. The upper plot in Figure 1 shows simulated quarterly GDP growth rates Y_t (black line) over a 25 year period with the hidden or unobserved level μ_{S_t} (grey horizontal lines) superimposed. The times when S_t changes state are indicated by the vertical grey lines so that S_t is initially in the first state and then cycles through $S_t = 2$, $S_t = 1$, $S_t = 2$ and finally ends up in the first state at the end of the series. Note the higher volatility in the second state $S_t = 2$. The lower plot shows the GDP series (black line) that is obtained by integrating the growth rates Y_t and, as before, the vertical grey lines indicate the times when the state changes.

This conceptually simple model is more versatile and more general than it might seem at first sight. In addition to allowing for switching level and volatility regimes as well as structural breaks in these parameters (see Kim and Nelson (1999b) for example), the deviations $\sigma_{S_t}X_t$ can also model non–Gaussian behaviour such as heavy tails using Gaussian mixture distributions. The latter follows from a judicious choice of parameters for the hidden Markov chain. Thus the model can be organised to be robust to outliers and other heavy tailed phenomena which is useful when analysing volatile data. In addition, through the ordered pairs (μ_j,σ_j) ($j = 1, \ldots, N$), the model provides a useful tool for investigating any relationships between the mean levels μ_j and the volatility levels σ_j of the states.

2.1 Specification

Given the length of the quarterly GDP series under study (92 observations) and the need for parsimonious models, we consider only the simple case where X_t (the original series of GDP growth rates Y_t corrected for level μ_{S_t} and volatility σ_{S_t}) is an $AR(1)$ process. Thus X_t satisfies

$$X_t = \rho X_{t-1} + \epsilon_t \qquad (|\rho| < 1; t = 0, \pm 1, \ldots) \tag{2}$$

where ϵ_t is Gaussian white noise with variance one. The latter condition serves to identify σ_t. However the procedures that we advocate are not restricted to this assumption. If sufficient quality data are available, then other models for X_t (e.g. $ARMA(p,q)$) and other distributions for ϵ_t can be fitted using a straightforward generalisation of the techniques described here.

The stationary finite Markov chain S_t is assumed to be ergodic and irreducible, but could otherwise have quite general structure. In particular S_t will be specified by its stationary transition probablities

$$P_{ij} = P(S_{t+1} = j|S_t = i) \qquad (i, j = 1, \ldots, N) \tag{3}$$

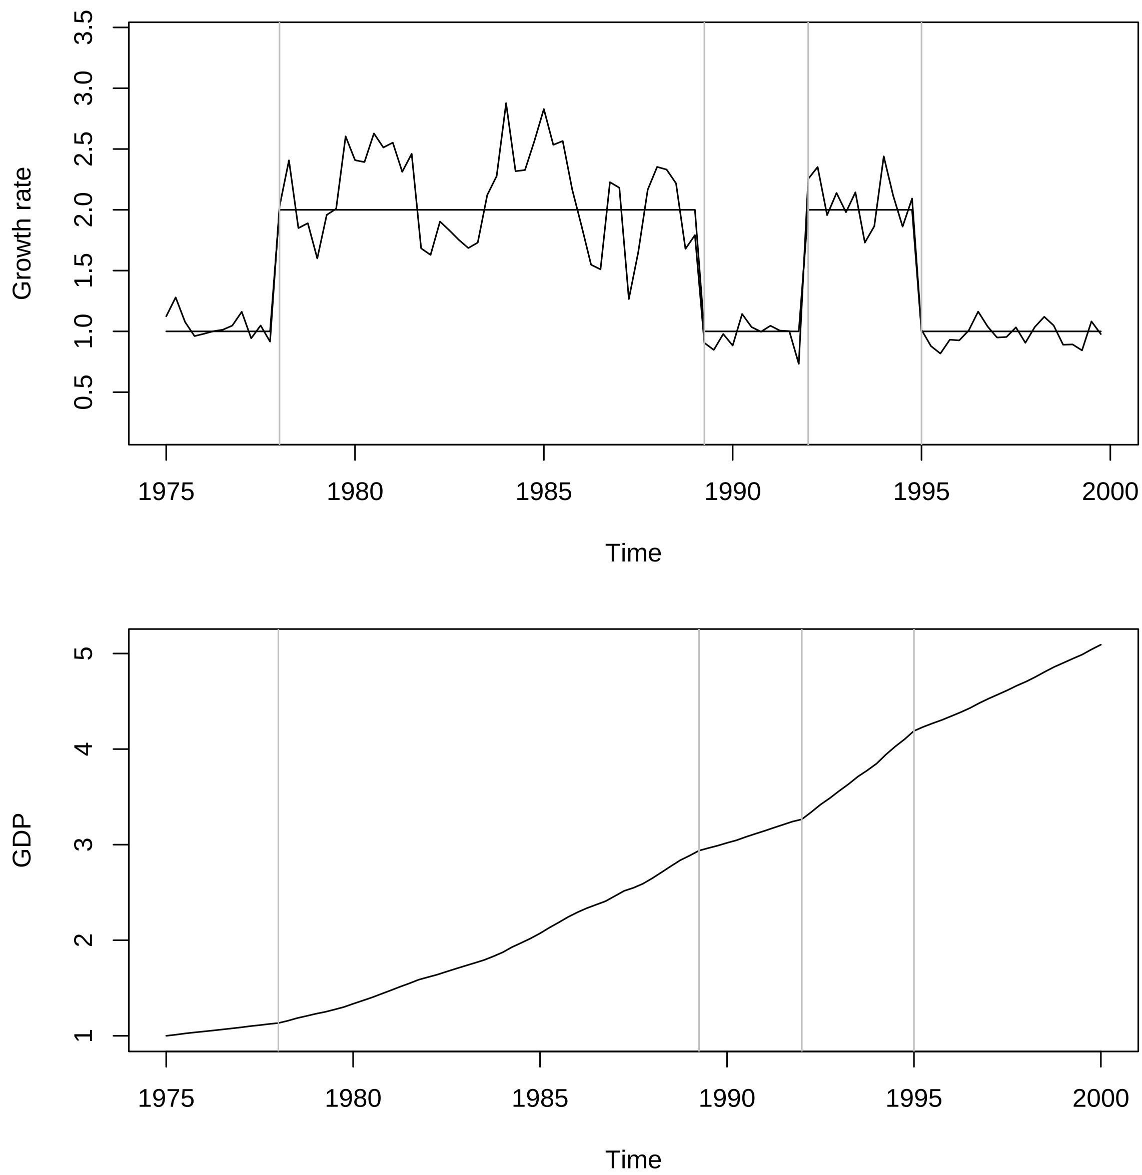

Figure 1: The upper plot shows simulated quarterly GDP growth rates (black line) with the unobserved mean level μ_{S_t} (grey horizontal lines) superimposed. The times when S_t changes state are indicated by the vertical grey lines. The lower plot shows the GDP series (black line) that is obtained by integrating the growth rates Y_t and the times when the state changes (vertical grey lines).

where the constraints $\sum_{j=1}^{N} P_{ij} = 1$ ($i = 1, \ldots, N$) imply a total of $N(N-1)$ parameters in all. For $N = 4$ states, for example, this would lead in principle to 12 parameters that would need to be estimated. The fact that the number of parameters required to specify S_t increases quadratically with N is a major weakness of the model. In practice N must be kept small or other simpler switching models adopted. We adopt both strategies in what follows.

Motivated by the need for more parsiminious models for S_t, we follow in the footsteps of McConnell and Perez-Quiros (2000) and consider more specific generating mechanisms. In their case S_t is specified by two independent Markov chains C_t and V_t which each take on the values 0 and 1. The chain C_t is intended to describe the growth regimes of the business cycle (recession when $C_t = 0$; growth when $C_t = 1$) with

$$\begin{aligned} P(C_t = 0) &= p_0 \\ P(C_t = 1) &= p_1 = 1 - p_0 \\ P(C_t = j|C_{t-1} = i) &= p_{ij} \qquad (i, j = 0, 1) \end{aligned} \tag{4}$$

and the transition probabilities satisfy $p_{01} = 1 - p_{00}$, $p_{10} = 1 - p_{11}$. Similarly V_t is intended to describe the stages of the volatility cycle (low when $V_t = 0$; high when $V_t = 1$) with

$$\begin{aligned} P(V_t = 0) &= q_0 \\ P(V_t = 1) &= q_1 = 1 - q_0 \\ P(V_t = j|V_{t-1} = i) &= q_{ij} \qquad (i, j = 0, 1) \end{aligned} \tag{5}$$

and $q_{01} = 1 - q_{00}$, $q_{10} = 1 - q_{11}$. In each case the chains are stationary so that the p_i and q_i must satisfy

$$p_0 = \frac{1 - p_{11}}{2 - p_{00} - p_{11}}, \qquad q_0 = \frac{1 - q_{11}}{2 - q_{00} - q_{11}}. \tag{6}$$

Given this structure, S_t takes on the values $1, \ldots, 4$ which represent the various combinations of the two cycles. More specifically we have the 1–1 mapping given in Table 1 where $S_t = 1 + V_t + 2C_t$. Columns 4 and 5 of Table 1 also illustrate a mapping of states to the business and volatility cycles. In what follows we shall use the word *regime* to denote a particular phase of the cycle (e.g. high growth phase of the business cycle) and define regimes to be suitable collections of states. Thus, in Table 1, the low growth phase of the business cycle is a regime

S_t	C_t	V_t	Business cycle	Volatility cycle	μ_{S_t}	σ_{S_t}
1	0	0	Low	Low	μ_1	σ_1
2	0	1	Low	High	μ_2	σ_2
3	1	0	High	Low	μ_3	σ_3
4	1	1	High	High	μ_4	σ_4

Table 1: 1–1 mapping of the state labels for S_t to those for C_t and V_t.

with two states ($S_t = 1$, $S_t = 2$) and the high phase of the volatility cycle is a regime with two states ($S_t = 2$, $S_t = 4$). Given this definition we now typically have a hierarchy of time scales with longer time–scale regimes comprising shorter time-scale states which, in turn, model the time series in the original time scale of the observations. Such a hierarchical classification of time scales is one of the features of hidden Markov models and provides a relatively simple and open structure on which to build an overall model for Y_t.

McConnell and Perez-Quiros (2000) consider the situation where the levels of the business cycle change whenever the volatility changes, but the levels of the volatility cycle are invariant to changes in the level of the business cycle. In other words, the business cycle is a function of volatility, but not vice versa. In this case there are only two distinct values for σ_{S_t} ($\sigma_1 = \sigma_3, \sigma_2 = \sigma_4$), but four distinct values for μ_{S_t} ($\mu_1, \mu_2, \mu_3, \mu_4$) corresponding to the two levels of each business cycle regime within the two levels of each volatility regime. An example of their model is given in Figure 2 which shows simulated US growth rates using the parameters fitted by McConnell and Perez-Quiros (2000) to actual US real GDP quarterly growth rates over the period 1953:2 to 1999:2. The sample path of this particular realisation illustrates the sustained periods of high growth and short periods of recession that we would expect for the US data. More generally, the ability to directly model the persistence of the cycles is a feature and potential strength of the Markov switching models.

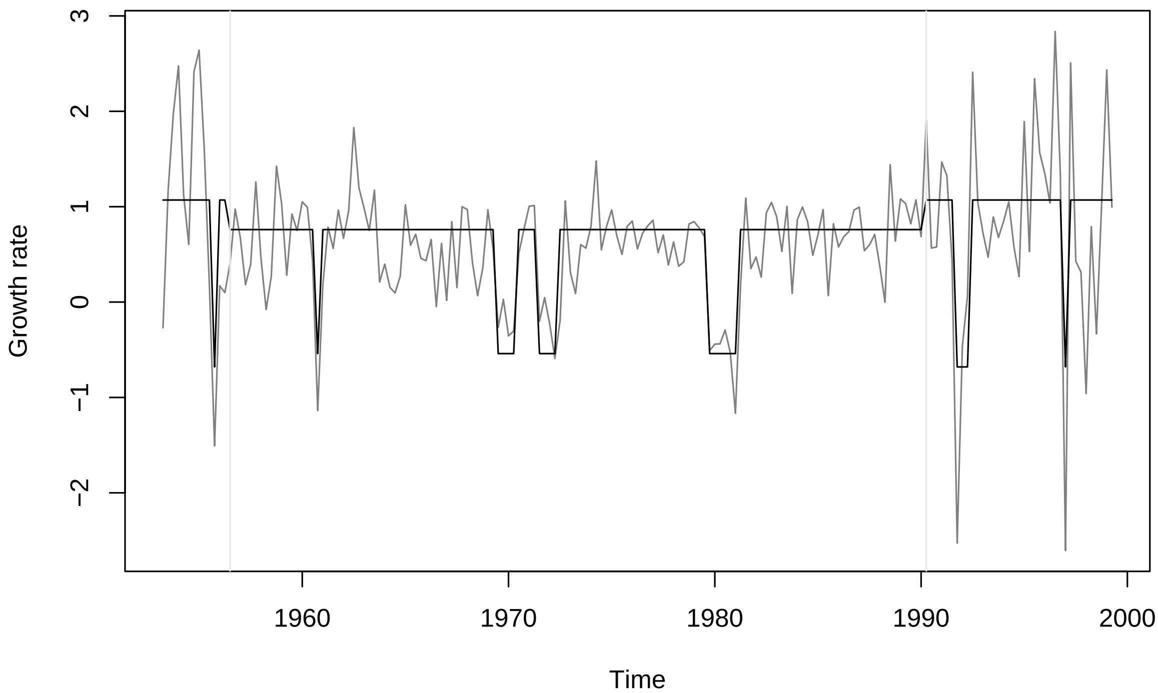

Figure 2: Simulated US growth rates using the parameters fitted by McConnell and Perez-Quiros (2000). The horizontal lines denote the levels of the business cycle (the two lowest levels are very similar and are associated with the low growth regime; the two highest with the high growth regime) and the vertical lines denote the time points when the volatility changes (the middle period is associated with the low volatility regime; the end periods with the high volatility regime).

For this study we adopt a more general view of the model and, unlike McConnell and Perez-Quiros (2000), do not necessarily impose any a priori restrictions on the σ_j ($j = 1, \ldots, 4$) or, for that matter, the μ_j ($j = 1, \ldots, 4$). Our starting point is the basic model described above with the full 13 parameters p_{jj}, q_{jj} (j = 0,1), μ_j, σ_j (j = 1, ..., 4) and ρ. More importantly, we shall consider a broader range of interpretations for the states $S_t = 1, \ldots, 4$ and will not necessarily be restricted to those implied by the 4th and 5th columns of Table 1. In this way we can use our parsimonious Markovian switching model as an approximation to a more general Markov chain S_t. Thus our model can be viewed as approximating a system with $N = 4$ states and transition probabilities (3) that potentially involve 12 free parameters, by a low–dimension system with 4 free parameters. Of course this approach could be extended further for larger values of N with even more parsimonious results. Such a strategy seems difficult to avoid given the relatively short times series under study.

In common with other disciplines where hidden Markovian models are used to good effect, the classification of states to regimes or, equivalently, the assigning of economic labels to states, is essentially a subjective process. It provides the economic analyst with an opportunity to interact with the model by vesting the regimes with meaning and interpretation useful for economic and policy analysis. In some situations this may be regarded as a potential weakness, but here we regard it as a major strength. The appropriate attribution of economic labels to states is an important aspect of the model fitting process which, in this case, is enhanced by the conceptually simple structure of the model.

We note in passing that the structural form adopted for the model (1) is not quite the same as that proposed by McConnell and Perez-Quiros (2000) and Kim and Nelson (1999) for example. The equivalent of their model in the case of $AR(1)$ errors is given by

$$Y_t = \mu_{S_t} + \rho(Y_{t-1} - \mu_{S_{t-1}}) + \sigma_{S_t}\epsilon_t \qquad (t = 0, \pm 1, \ldots)$$

or, equivalently,

$$\frac{Y_t - \mu_{S_t}}{\sigma_{S_t}} = \rho \frac{\sigma_{S_{t-1}}}{\sigma_{S_t}} \frac{Y_{t-1} - \mu_{S_{t-1}}}{\sigma_{S_{t-1}}} + \epsilon_t.$$

The latter model is almost identical to (1) and (2) except at the times when the volatility state changes. In particular, note that the correlation between Y_t and Y_{t-1}, given S_t and S_{t-1}, is not constant as in (1), but time varying.

The model (1) is an example of an HMM (Hidden Markov Model) first proposed by Baum and Petrie (1966). General references to HMM modelling include Levinson, Rabiner and Sondhi (1983), Rabiner (1989), Elliot, Aggoun and Moore (1995) and MacDonald and Zucchini (1997). Following the lead of Goldfeld and Quandt (1973), and Hamilton (1989), these and related methods have been used widely in economic contexts (see Engle and Hamilton (1990), Hamilton and Susmel (1994), Kim (1994), Kim and Nelson (1999a, 1999b), McConnell and Perez–Quiros (2000), Kontolemis (2001) for example). In particular Krolzig (1997) provides a comprehensive and thorough account of the theory and inference for Markov switching vector autoregressions with application to business cycle analysis.

2.2 Examples

The full model adopted encompasses many other reduced models of interest. These include the following.

AR(1) Setting

$$\mu_1 = \mu_2 = \mu_3 = \mu_4, \quad \sigma_1 = \sigma_2 = \sigma_3 = \sigma_4$$

and $p_{00} = 1$, $p_{11} = 0$, $q_{00} = 1$, $q_{11} = 0$ yields a simple AR(1) model with constant mean and volatility. This represents a null model with no cycles present.

Hamilton The seminal model proposed by Hamilton (1989) is obtained by setting

$$\mu_1 = \mu_2, \quad \mu_3 = \mu_4, \quad \sigma_1 = \sigma_2 = \sigma_3 = \sigma_4$$

and $q_{00} = 1$, $q_{11} = 0$. For the Hamilton model the level μ_{S_t} switches between two values, the volatility σ_{S_t} is constant, and V_t is constrained to be 0. The total number of free parameters is 6.

MPQ As noted before, the model proposed by McConnell and Perez-Quiros (2000) is obtained by setting

$$\sigma_1 = \sigma_3, \quad \sigma_2 = \sigma_4$$

so that the level μ_{S_t} switches between four values and the volatility σ_{S_t} switches between two values. Alternatively μ_{S_t} can be regarded as switching between two basic levels (high and low say) which, in turn, are dependent on which of the two volatility states the process is in. Here the total number of free parameters is 11.

Hamilton with outliers A simple variant of the Hamilton model that allows for outliers is obtained by setting

$$\mu_1 = \mu_2, \quad \mu_3 = \mu_4, \quad \sigma_1 = \sigma_3, \quad \sigma_2 = \sigma_4 = 3\sigma_1$$

and $q_{00} = 0.99 = 1 - q_{11}$. This assumes, somewhat arbitrarily, that outliers occur independently about 1% of the time and, when they do ($V_t = 1$), they are drawn from a Gaussian distribution with large standard deviation. Given that outliers are likely to occur infrequently, such assumptions offer a simple way to build models that are resistant to outliers. Like the Hamilton model, this model has 6 free parameters.

Hamilton with non–Gaussian errors Non–Gaussian errors can be accommodated within the Hamilton model by setting

$$\mu_1 = \mu_2, \quad \mu_3 = \mu_4, \quad \sigma_1 = \sigma_3, \quad \sigma_2 = \sigma_4$$

and $q_{00} = 1 - q_{11}$. The last condition ensures that the V_t are independent Bernoulli random variables with $P(V_t = 0) = q_{00}$ and $P(V_t = 1) = q_{11}$. Then the (marginal) distribution of the errors $\sigma_{S_t} X_t$ is a mixture of Gaussian distributions which can be chosen to mimic some other distribution, such as a heavy tailed distribution. This allows the model some flexibility to be robust to distributional assumptions. Here the model has 7 free parameters.

Other Many other reduced models are possible. Such models are referred to informally in Section 4 as m–v models where the numbers m and v refer to the number of distinct mean parameters μ_j and volatility parameters σ_j respectively. For example setting

$$\mu_1 = \mu_2, \quad \mu_3 = \mu_4, \quad \sigma_1 = \sigma_2, \quad \sigma_3 = \sigma_4$$

and $q_{00} = 1$, $q_{11} = 0$ is an example of a 2–2 model where the business and volatility cycles coincide and each phase of the business cycle has its own mean and standard deviation. Setting

$$\sigma_1 = \sigma_2 = \sigma_3 = \sigma_4$$

and retaining 4 levels for the μ_j is an example of a 4–1 model. In the latter case the volatility is constant and the 4 states can be allocated to two or more business cycle regimes.

2.3 Second order properties

Here we determine the mean and autocovariance function of the stationary time series Y_t where Y_t follows (1), X_t follows (2) and the hidden Markov chain S_t has transition probability matrix $\mathbf{P} = \{P_{ij}\}$ given by (3). The stationary distribution of the chain is given by the row vector $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_N)$ where $\boldsymbol{\pi}\mathbf{P} = \boldsymbol{\pi}$ and $P(S_t = j) = \pi_j$ ($j = 1, \ldots, N$) with $\sum_{j=1}^{N} \pi_j = 1$. Given this general structure, we now determine the mean, autocovariance and autocorrelation functions of Y_t.

First write

$$Y_t = L_t + D_t \qquad (t = 0, \pm 1, \ldots)$$

where $L_t = \mu_{S_t}$ denotes the level of Y_t and the $D_t = \sigma_{S_t} X_t$ denote the deviations of Y_t from L_t. The level L_t has mean

$$\mu_L = E(\mu_{S_t}) = \sum_{j=1}^{N} \mu_j \pi_j$$

and autocovariance function

$$\gamma_L(t) = E\{\mu_{S_s}\mu_{S_{s+t}}\} - \mu_L^2 = \sum_{j=1}^{N}\sum_{k=1}^{N} \mu_j \mu_k (P_{jk}^{|t|}\pi_j - \pi_j \pi_k) \qquad (t = 0, \pm 1, \ldots)$$

where $P_{jk}^n = P(S_{s+n} = k | S_s = j)$ ($n \geq 0$) is the typical element of the matrix $\mathbf{P}^n$. In particular the variance of L_t is given by

$$\gamma_L(0) = \sum_{j=1}^{N} (\mu_j - \mu_L)^2 \pi_j.$$

The deviations D_t have mean zero and autocovariance function

$$\gamma_D(t) = \gamma_X(t) E\{\sigma_{S_s}\sigma_{S_{s+t}}\} = \gamma_X(t) \sum_{j=1}^{N}\sum_{k=1}^{N} \sigma_j \sigma_k P_{jk}^{|t|} \pi_j \qquad (t = 0, \pm 1, \ldots)$$

where

$$\gamma_X(t) = \frac{\rho^{|t|}}{1 - \rho^2} \qquad (t = 0, \pm 1, \ldots)$$

is the autocovariance function of X_t and

$$\gamma_D(0) = \frac{\sum_{j=1}^N \sigma_j^2 \pi_j}{(1-\rho^2)}$$

is the variance of D_t.

Noting that L_t and D_t are mutually uncorrelated, the mean and autocovariance function of $Y_t = L_t + D_t$ are now given by

$$E(Y_t) = \mu_L, \qquad \gamma_Y(t) = \gamma_L(t) + \gamma_D(t) \qquad (t = 0, \pm 1, \ldots) \tag{7}$$

respectively. Furthermore, the autocorrelation function of Y_t is given by

$$\rho_Y(t) = \alpha \rho_L(t) + (1-\alpha)\rho_D(t) \qquad (0 \leq \alpha \leq 1; t = 0, \pm 1, \ldots) \tag{8}$$

where $\rho_L(t) = \gamma_L(t)/\gamma_L(0)$, $\rho_D(t) = \gamma_D(t)/\gamma_D(0)$ are the autocorrelation functions of L_t, D_t respectively, $\alpha = \delta/(1+\delta)$ and the *signal to noise ratio* δ is given by

$$\delta = \frac{\gamma_L(0)}{\gamma_D(0)} = (1-\rho^2)\frac{\sum_{j=1}^N (\mu_j - \mu_L)^2 \pi_j}{\sum_{j=1}^N \sigma_j^2 \pi_j}.$$

The larger the absolute size of the deviations of the levels μ_j from the overall mean $E(Y_t) = \mu_L$ relative to the size of the standard deviations σ_j, the larger δ and the closer α is to one and vice–versa. Since $\rho_D(t)$ is dominated by the autoregressive autocorrelation function $\rho^{|t|}$ and $\rho_L(t)$ typically decays much more slowly to zero for the applications we have in mind, the values of δ and α play a key role in determining the nature of $\rho_Y(t)$. In practice we will tend to see the geometric correlation structure $\rho^{|t|}$ when δ is small, and the longer memory autocorrelation structure of L_t whan δ is large. Indeed, using observed correlations alone, it would be very difficult to extract the volatility and autoregressive correlation structure from the deviations D_t when δ is large, and difficult to determine the autocorrelation structure of the levels L_t when δ is small.

Finally we note that $\gamma_L(t)$ and $\gamma_D(t)$ are both linear combinations of geometrically decaying terms. This leads to the observation that L_t and D_t have the covariance structure of $ARMA(p,q)$ processes with $p \leq N-1$, $q \leq N-2$ and $p \leq N$, $q \leq N-1$ respectively. Thus Y_t is a second–order stationary $ARMA(p,q)$ process with $p \leq 2N-1$, $q \leq 2N-2$.

3 Fitting the model

Given observations $Y_1, \ldots, Y_T$ our general strategy is to fit the model (1) using maximum likelihood and the EM algorithm (Dempster, Laird and Rubin (1977)) with the choice of model orders guided by the BIC criterion. The latter selects the model that minimises

$$BIC = -2 \text{ log likelihood} + p \ln T$$

with respect to the model order p. As in the case of AIC, this criterion trades model fit against model complexity. The EM algorithm can be used to obtain exact maximum likelihood estimates for certain models. However, in almost all cases we use it to explore the likelihood surface and obtain approximate maximum likelihood estimates which, in turn, are further refined using direct maximum likelihood. In the latter case we take advantage of the EM algorithm's relative insensitivity to choice of initial values. Issues such as the determination of the standard errors of the parameters and the extraction of the trend $E(Y_t|S_t) = \mu_{S_t}$ and volatility $\mathrm{Var}(Y_t|S_t) = \sigma^2_{S_t}\mathrm{Var}(X_t) = \sigma^2_{S_t}/(1-\rho^2)$ from the data will also be considered.

Given $\mathbf{S} = (S_1, \ldots, S_T)$ the density of $\mathbf{Y} = (Y_1, \ldots, Y_T)$ is given by

$$f(\mathbf{Y}|\mathbf{S}) = \frac{\sqrt{1-\rho^2}}{\sigma_{S_1}\sqrt{2\pi}} e^{-\frac{1}{2}(1-\rho^2)X_1^2} \prod_{t=2}^{T} \frac{1}{\sigma_{S_t}\sqrt{2\pi}} e^{-\frac{1}{2}(X_t - \rho X_{t-1})^2} \tag{9}$$

where

$$X_t = \frac{Y_t - \mu_{S_t}}{\sigma_{S_t}} \qquad (t = 1, \ldots, T).$$

The density of the S_t, or equivalently the C_t, V_t, is given by

$$f(\mathbf{S}) = \pi_{S_1} \prod_{t=2}^{T} P_{S_{t-1}S_t} = \left(p_{C_1} \prod_{t=2}^{T} p_{C_{t-1}C_t} \right) \left(q_{V_1} \prod_{t=2}^{T} q_{V_{t-1}V_t} \right) \tag{10}$$

where the π_j, P_{ij} are as defined in Section 2.3 and are functions of the 4 parameters p_{00}, p_{11}, q_{00}, q_{11}. Thus the log–likelihood of $\mathbf{Y}$ and $\mathbf{S}$ is given by

$$\log L_c(\boldsymbol{\theta}) = \log f(\mathbf{Y}|\mathbf{S}) + \log f(\mathbf{S}) \tag{11}$$

where

$$\begin{aligned} \log f(\mathbf{Y}|\mathbf{S}) &= -\tfrac{1}{2}T\log(2\pi) + \tfrac{1}{2}\log(1-\rho^2) - \tfrac{1}{2}(1-\rho^2)X_1^2 \\ &\quad -\tfrac{1}{2}\sum_{t=1}^{T}\log\sigma^2_{S_t} - \tfrac{1}{2}\sum_{t=2}^{T}(X_t - \rho X_{t-1})^2 \\ \log f(\mathbf{S}) &= \log \pi_{S_1} + \sum_{t=2}^{T} \log P_{S_{t-1}S_t}. \end{aligned}$$

The vector $\boldsymbol{\theta}$ in (11) denotes the model parameters p_{jj}, q_{jj} ($j = 0, 1$), μ_j, σ_j ($j = 1, \ldots, 4$) and ρ so that $\boldsymbol{\theta}$ has dimension 13. In keeping with EM terminology we call $\log L_c(\boldsymbol{\theta})$ the log–likelihood of the *complete data* $\mathbf{Y}, \mathbf{S}$.

However it is the likelihood of $\mathbf{Y}$ (the *incomplete data*) that we must determine since this is the only data we have available. The likelihood of $\mathbf{Y}$ is given by

$$L(\boldsymbol{\theta}) = \sum_{\mathbf{S}} f(\mathbf{Y}|\mathbf{S}) f(\mathbf{S}) \tag{12}$$

where $\sum_{\mathbf{S}}$ is over all possible realisations of $\mathbf{S}$. It is $L(\boldsymbol{\theta})$ or $\log L(\boldsymbol{\theta})$ that should ideally be optimised with respect to $\boldsymbol{\theta}$ to determine the maximum likelihood estimator $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$. The more complicated structure of $\log L(\boldsymbol{\theta})$ makes it a more difficult function to directly optimise by comparison to $\log L_c(\boldsymbol{\theta})$. This and other reasons lead us to first consider the EM algorithm.

If the states S_t were known, then it is the relatively simple complete log–likelihood $\log L_c(\boldsymbol{\theta})$ that would be optimised to determine estimators of $\boldsymbol{\theta}$. Given only the observations $\mathbf{Y}$ the best (quadratic loss) predictor of $\log L_c(\boldsymbol{\theta})$ is

$$Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0) = E(\log L_c(\boldsymbol{\theta})|\mathbf{Y}) \tag{13}$$

where the expectation operator E is with respect to the true distribution indexed by $\boldsymbol{\theta}_0$. Given an initial estimate of $\boldsymbol{\theta}_0$ a new estimate can be found by maximising $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$ with respect to the parameters $\boldsymbol{\theta}$. The new estimate can, in turn, be used for $\boldsymbol{\theta}_0$ and so on. This recursion forms the basis of the celebrated EM algorithm (Dempster, Laird and Rubin (1977)) where the determination of the conditional expectation $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$ is referred to as the E–step and its maximisation with respect to $\boldsymbol{\theta}$ the M–step. Under certain general conditions it can be shown that the sequence of estimates constructed in this way yields monotonically increasing values of $L(\boldsymbol{\theta})$ and converges to the maximum likelihood estimator $\hat{\boldsymbol{\theta}}$ for the incomplete data. Thus the EM algorithm provides an alternative method of maximising the log–likelihood $\log L(\boldsymbol{\theta})$.

The computational efficiency of the EM algorithm is greatly enhanced if the E and M steps are readily evaluated, particularly the M step where simple closed form solutions are desired. In this case the algorithm is particularly easy to implement. In practice the EM algorithm is often more robust to the choice of initial starting values than direct maximum likelihood which, if numerical optimisation procedures are used, tends to converge to a local rather than a global maximum. However, although better at identifying the region containing the global maximum, the EM algorithm can often be slow to converge in the vicinity of the global maximum. One reason for this is that the EM criterion $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$ is essentially a smoothed form of a log–likelihood and so the algorithm is less likely to converge to a local maximum than direct maximum likelihood, but more likely to converge slowly near the maximum due to a flattened log–likelihood surface. These observations and design objectives underpin the development that follows.

From (11) and (13) we obtain

$$\begin{aligned} Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0) \;=\; & -\tfrac{1}{2}T\log(2\pi) + \tfrac{1}{2}\log(1-\rho^2) - \tfrac{1}{2}(1-\rho^2)\sum_{j=1}^{4}\left(\frac{Y_1-\mu_j}{\sigma_j}\right)^2\gamma_1(j) \\ & -\tfrac{1}{2}\sum_{j=1}^{4}\log\sigma_j^2\sum_{t=1}^{T}\gamma_t(j) - \tfrac{1}{2}\sum_{j=1}^{4}\sum_{k=1}^{4}\sum_{t=1}^{T-1}\left(\frac{Y_{t+1}-\mu_k}{\sigma_k} - \rho\frac{Y_t-\mu_j}{\sigma_j}\right)^2\gamma_t(j,k) \\ & +\sum_{j=1}^{4}(\log\pi_j)\gamma_1(j) + \sum_{j=1}^{4}\sum_{k=1}^{4}\log P_{jk}\sum_{t=1}^{T-1}\gamma_t(j,k) \end{aligned}$$

where

$$\gamma_t(j,k) = P(S_t=j, S_{t+1}=k|\mathbf{Y},\boldsymbol{\theta}_0), \qquad \gamma_t(j) = P(S_t=j|\mathbf{Y},\boldsymbol{\theta}_0) = \sum_{k=1}^{4}\gamma_t(j,k).$$

The probabilities $\gamma_t(j)$ and $\gamma_t(j,k)$ are functions only of the initial parameters $\boldsymbol{\theta}_0$, the data $\mathbf{Y}$, but not the parameters $\boldsymbol{\theta}$. They need to be determined prior to evaluating and optimising $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$.

Efficient recursive algorithms are given in the Appendix for evaluating the $\gamma_t(j)$ and the $\gamma_t(j,k)$. An important by–product of these recursions is the evaluation of the exact likelihood $L(\boldsymbol{\theta})$ given by (12). Thus we now have an appropriate computational framework in place for calculating maximum likelihood estimates by direct maximum likelihood (using numerical optimisation routines) as well as by the EM algorithm.

However the $\gamma_t(j)$ and $\gamma_t(j,k)$ are also useful in their own right to extract estimates of stochastic parameters such as the trend μ_{S_t} and volatility σ_{S_t}. For example the best (quadratic loss) estimate of μ_{S_t} given the data $\mathbf{Y}$ is

$$E(\mu_{S_t}|\mathbf{Y}) = \sum_{j=1}^{4} \mu_j \gamma_t(j) \tag{14}$$

and the best (quadratic loss) estimate of $\mathrm{Var}(Y_t|S_t)$ given the data is

$$E(\frac{\sigma_{S_t}^2}{1-\rho^2}|\mathbf{Y}) = \frac{1}{1-\rho^2} \sum_{j=1}^{4} \sigma_j^2 \gamma_t(j). \tag{15}$$

These estimates of the time varying mean and variance of Y_t are used as informal diagnostic graphical measures in the applications sections. Equally importantly, the $\gamma_t(j)$ and $\gamma_t(j,k)$ also provide useful measures for identifying and classifying the most likely regimes for the hidden cycles.

Despite the relatively simple structure of $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$ as a function of $\boldsymbol{\theta}$, analytic solutions for the value of $\boldsymbol{\theta}$ that maximises $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$ will only exist in certain situations and then only if certain approximations are made. An important case is where $\rho = 0$ and all other parameters are distinct. Then, retaining only those terms of order T in $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$, the estimates of the parameters that maximise $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$ are given by

$$\tilde{\mu}_j = \frac{\sum_{t=1}^{T} \gamma_t(j) Y_t}{\sum_{t=1}^{T} \gamma_t(j)}, \qquad \tilde{\sigma}_j^2 = \frac{\sum_{t=1}^{T} \gamma_t(j)(Y_t - \tilde{\mu}_j)^2}{\sum_{t=1}^{T} \gamma_t(j)} \qquad (j = 1, \ldots, 4) \tag{16}$$

and

$$\begin{aligned} \tilde{p}_{00} &= \frac{\sum_{t=1}^{T-1} P(C_t = 0, C_{t+1} = 0|\mathbf{Y}, \boldsymbol{\theta}_0)}{\sum_{t=1}^{T-1} P(C_t = 0|\mathbf{Y}, \boldsymbol{\theta}_0)} = \frac{\sum_{t=1}^{T-1} \sum_{j=1}^{2} \sum_{k=1}^{2} \gamma_t(j,k)}{\sum_{t=1}^{T-1} \sum_{j=1}^{2} \gamma_t(j)} \\ \tilde{p}_{11} &= \frac{\sum_{t=1}^{T-1} P(C_t = 1, C_{t+1} = 1|\mathbf{Y}, \boldsymbol{\theta}_0)}{\sum_{t=1}^{T-1} P(C_t = 1|\mathbf{Y}, \boldsymbol{\theta}_0)} = \frac{\sum_{t=1}^{T-1} \sum_{j=3}^{4} \sum_{k=3}^{4} \gamma_t(j,k)}{\sum_{t=1}^{T-1} \sum_{j=3}^{4} \gamma_t(j)} \end{aligned} \tag{17}$$

with the analogous expressions for $\tilde{q}_{00}$, $\tilde{q}_{11}$ involving V_t instead of C_t. Equations (16), (17) provide the required EM recursions which will converge to the maximum likelihood estimate of the parameters in this case where ρ is constrained to be zero and the μ_j, σ_k are distinct. Although there are other cases where analytic EM recursions can be found, this particular case was used to explore the log–likelihood surface to identify suitable initial estimates for direct maximum likelihood using numerical optimisation procedures.

Table 2 provides a summary of the fitting procedure adopted in the applications given in the following sections. Using these methods and strategies, we now fit the various models considered to New Zealand GDP data.

1. Use the EM recursions (16), (17) to explore the log–likelihood surface. and obtain a range of suitable initial estimates for the maximum likelihood estimate $\hat{\boldsymbol{\theta}}$

2. Starting from these initial estimates, use numerical optimisation procedures to directly maximise the log–likelihood $\log L(\boldsymbol{\theta})$ subject to parameter constraints ($|\rho| < 1$, $0 \leq p_{jj} \leq 1$, $0 \leq q_{jj} \leq 1$ for $j = 0, 1$, and $\sigma_j > 0$ for $j = 1, \ldots, 4$). Here the exact log–likelihood is evaluated using the recursions given in the Appendix.

3. Determine the standard errors of the maximum likelihood estimates from the information matrix obtained from the Hessian provided by the optimisation procedure.

4. Examine the resulting estimates, BIC values etc and suitable graphical diagnostics to assess goodness of fit.

5. Identify and classify the most likely regimes for hidden business and. volatility cycles

Table 2: Summary of fitting procedure.

4 HMM models for New Zealand GDP growth

This section identifies shifts in mean growth rates and volatilities by fitting Markov switching models to growth rates for total GDP and five production sectors that make up total GDP. The five sectors are *Services*, *Government and Community Services*, *Manufacturing*, *Primary*, and *Construction* as defined in Table 3. This work complements and builds on Buckle, Haugh and Thomson (2001) which attempts to identify the evolution of local means and volatility of quarterly growth rates for New Zealand real GDP and its production sectors using simple moving average techniques. That paper also decomposes aggregate GDP growth and its volatility into contributions from the individual sectors.

The GDP series used in this paper are quarterly real seasonally adjusted chain–linked production GDP for the period 1978:1 to 2000:4. The series are Statistics New Zealand new official quarterly chain series from 1987:2 onwards appended to a calibrated chain series for the period back to 1978:1. The calibration procedure is explained in Haugh (2001) and the same GDP series are used in Buckle, Haugh and Thomson (2001). The calibration procedure exploits the statistical relationship between the period of overlapping official chain–linked and ex–official fixed–weight series (1987:2 to 2000:2) which is then used to derive series for each production sector and for total real GDP for the period from 1978:1 to 1987:1. These calibrated series are intended to approximate the chain–linked series over this period and are combined with the respective 1987:2 to 2000:4 chain–linked series available from Statistics New Zealand to form consistent time series data for each sector over the period 1978:1 to 2000:4.

The choice of models to fit to GDP and its sectors was informed by the analysis of growth levels and volatility reported in Buckle, Haugh and Thomson (2001), including visual inspection of quarterly growth rates, and moving averages and standard deviations of GDP and each of

Sector name	Chain linked industries included in the sector
Services	Communications + Electricity, Gas & Water + Combined Wholesale Trade + Transport & Storage + Finance, Insurance, Business Services & Real Estate + Owner Occupied Dwellings
Government and Community Services	Personal and Community Services + Central Govt & Defence + Local Govt Services
Primary	Agriculture + Fishing + Forestry + Primary Food Manufacturing
Manufacturing	Textiles + Wood & Paper Products + Printing & Publishing + Petroleum etc + Non–Metallic Mineral Products Manufacturing + Basic Metals + Machinery & Equipment + Other Manufacturing + Other Food Manufacturing
Construction	Construction

Table 3: Industry composition of the five production sectors.

the sectors. Examination of the moving averages can be very useful in determining which local means a series appears to move around and the number of means to include in the HMM model. The moving standard deviations can be used similarly to determine the local standard deviations and how many volatility regimes there might be in the data. However, since the standard deviation is dependent on where the mean is placed it is not always as straightforward as it may seem. In other words, a change in the series may be interpreted as a shift in the local mean or a change in the standard deviation around a constant mean. The HMM is a tool that can be used to more fully understand whether various features of the data are shifts in local means or a change in volatility.

Visual inspection of GDP and sector quarterly growth rates suggest that the properties vary markedly across sectors and that allowing for different HMM models with varying numbers of states and varying means and standard deviations is appropriate. An initial model for each sector is selected for fitting based on prior analysis of means and standard deviations using centred moving average estimates of mean quarterly growth. These results are then used to inform any changes in the model being fitted. For example, if a four mean and two standard deviation model (4–2 model) is estimated, but two of the four fitted means are almost identical, a three mean and two standard deviation model (3–2 model) is fitted. This general to specific approach is supplemented by fitting simpler models with fewer parameters, such as the Hamilton two mean and single standard deviation model (2–1 model), to some sectors to obtain more robust estimates of the means which are then compared against the means estimated by more complicated models.

The AIC (Akaike Information Criterion) and BIC (Bayesian Information Criterion) model selection criteria were used to help select between competing models. This was supplemented by the criterion that the fitted model exhibit persistence in the sense that most regimes would be expected to last for a number of consecutive quarters before a switch takes place. An economy is unlikely to switch between high growth and low growth regimes every quarter because of the underlying economic process, which tends to show ongoing reinforcing behaviour that lasts more than one quarter. For example, in the high growth regime firms may be increasing investment, which leads to increased aggregate output and income, which in turn leads to more demand and so on. On this basis, the preferred model for a series that oscillates between extreme values every quarter for example, would be a constant mean with high volatility rather than two means at each extreme value even if the AIC and BIC favoured the latter model. Table 4 describes the preferred estimated HMM models for GDP and each production sector, and the parameter estimates for each of these models.

Data **Model**	**GDP** **Ham**	**GDP** **3–2**	**Ser** **MPQ**	**Gov** **2–2**	**Man** **Ham**	**Man** **MPQ**	**Man** **3–1**	**Pri** **3–2**	**Con** **3–1**
p_{00}	0.85	0.94	0.87	0.97	0.81	0.72	0.98	0.94	0.65
p_{11}	0.75	0.48	0.70	0.98	0.76	0.72	0.56	0.98	0.65
q_{00}	1.00	0.83	0.98	0.98	1.00	0.97	0.73	0.85	0.00
q_{11}	0.00	0.97	0.98	0.92	0.00	0.72	0.75	0.73	0.86
μ_1	0.15	1.23	0.68	0.28	-0.73	-1.01	-1.16	-1.74	-10.29
μ_2	0.15	0.25	0.04	0.28	-0.73	-3.07	1.59	4.88	-1.28
μ_3	1.27	2.06	1.25	0.81	1.89	1.74	4.76	0.97	2.66
μ_4	1.27	2.06	1.67	0.81	1.89	3.83	4.76	0.97	2.66
σ_1	0.98	0.22	0.66	0.53	1.77	1.23	1.45	1.89	3.98
σ_2	0.98	1.01	0.86	1.08	1.77	1.89	1.45	1.89	3.98
σ_3	0.98	0.22	0.66	0.53	1.77	1.23	1.45	3.30	3.98
σ_4	0.98	1.01	0.86	1.08	1.77	1.89	1.45	3.30	3.98
ρ	-0.03	-0.08	-0.26	-0.35	-0.14	-0.40	-0.29	-0.11	-0.26
$\log L$	-140.13	-135.02	-127.74	-95.01	-201.44	-197.66	-199.21	-241.38	-280.65
AIC	292.26	290.04	277.48	208.02	414.88	417.32	416.42	502.76	579.3
BIC	307.39	315.26	305.22	230.72	430.01	445.06	439.12	527.98	602
Parameters	6	10	11	9	6	11	9	10	9

Table 4: Parameter estimates for the HMM models fitted to GDP and sector growth rates. The sectors are Services (**Ser**), Government and Community Services (**Gov**), Primary (**Pri**), Manufacturing (**Man**) and Construction (**Con**) whose composition is given in Table 3. The models fitted are as indicated with **Ham** denoting the Hamilton model.

4.1 Aggregate real GDP

The Hamilton model, originally fitted to the US GNP growth rates, appears to successfully capture the dynamics of New Zealand real GDP. This model has also been successfully fitted to real GDP dynamics for several other countries (see for example Krolzig (1997)).

The top panel of Figure 3 shows the quarterly GDP series with the trend estimated from (14) and also from an 11–quarter triangular moving average for comparison. The second panel of Figure 3 plots the probability of being in the high growth regime of the business cycle. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the panel at the bottom of Figure 3. In particular the low–growth mean is estimated to be 0.15 percent per quarter and the high–growth mean is estimated to be 1.27 percent per quarter.

The Hamilton model indicates the New Zealand economy has experienced five upswings from low to high mean growth between 1978 and 2000, where the economy is regarded as being in

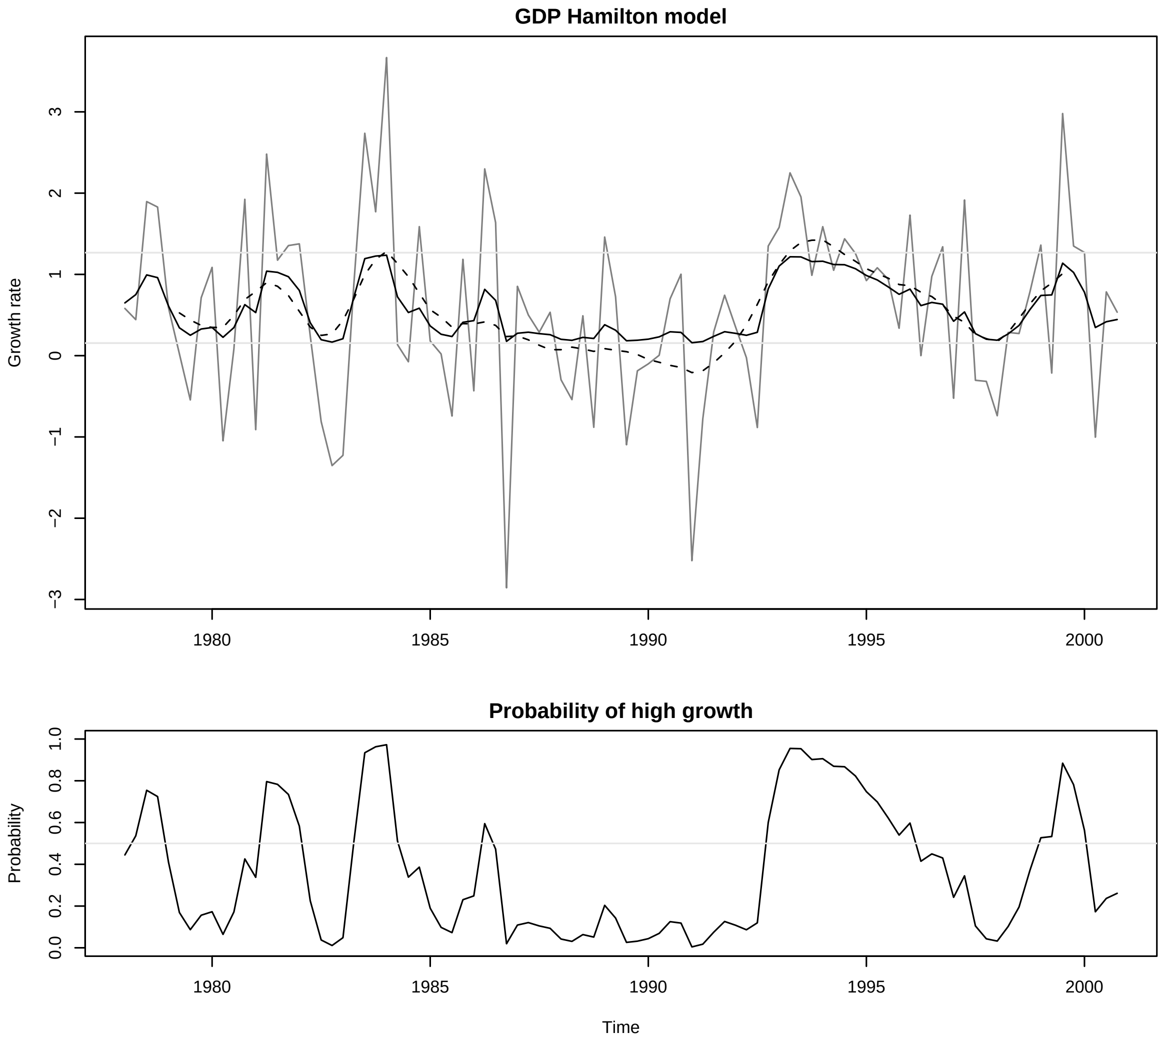

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1, 2	0.15	0.98	Low growth	Constant volatility
3, 4	1.27	0.98	High growth	Constant volatility

Figure 3: Results of fitting the Hamilton model to quarterly GDP growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

a high growth regime when the probability of being in that state is 50 percent or greater (otherwise it is defined as being in the low growth regime). According to this model, New Zealand experienced three switches from low to high growth regimes between 1978 and 1984 (these upswings are dated as follows: 1978:2–1978:4, 1981:2–1982:1, 1983:3–1984:2), a switch to a period of sustained high growth from 1992:3 to 1996:1, and another switch to the high growth regime at the end of the sample period (1999:1 to 2000:2). The Hamilton model also picks out 1986:2 as a period when GDP was in the high growth regime, but this was probably the effect of increased spending in anticipation of the introduction of GST on 1 October 1986. With the exception of this mid–1986 spike, the economy was in the low growth regime of the business cycle from 1984:3 to 1992:3.

Although a simple and parsimonious model, which is important when the number of observations (in this case 92) is not large, the Hamilton model seems nevertheless to be able to extract plausible business cycles from New Zealand GDP data that correspond to cycles derived from other trend measures, as discussed in Section 5.

Evidence of a decline in the standard deviation of New Zealand real GDP growth provided in Buckle, Haugh and Thomson (2001) suggests however that a richer HMM model with more than one standard deviation may be more appropriate. As a first step we fitted the MPQ model which has four mean growth rates and two standard deviations. This model was used by McConnell and Perez–Quiros (2000) to show evidence of breaks in US GDP volatility and allows the means in both growth regimes of the cycle to vary according to the level of volatility.

Fitting the MPQ model indicated two GDP volatility states in NZ real GDP growth, but only three distinct mean growth rates. Two of the estimated four mean growth rates (the two high means) were almost equal. On the basis of this evidence, a three means and two standard deviations model was fitted to NZ real GDP data (3–2 model).

In contrast to the Hamilton model which has two states (high and low mean growth states with a common standard deviation), the fitted HMM 3–2 model has four states. Of these, three (S_t = 1, 3, 4) are classified as belonging to the high growth regime and one (S_t = 2) is classified as belonging to the low growth regime. The classification of states to regimes is shown in the table at the bottom of Figure 4. The high growth regime has estimated mean growth rates of 1.23 percent per quarter and 2.06 percent per quarter. The latter picks out two short duration periods in 1984 and 1994 when quarterly real GDP growth rates were unusually high. The other high–growth mean and the low–growth mean are close to those for the Hamilton model.

Here the probability of a high growth regime occurring is $P(S_t = 1|\mathbf{Y}) + P(S_t = 3|\mathbf{Y}) + P(S_t = 4|\mathbf{Y})$ and this is plotted in the middle panel of Figure 4. This results in the identification of four switches from low growth to high growth regimes, one less than the number identified by the Hamilton model (excluding the 1986 GST spike). The periods of high growth regimes were as follows: 1981:3–1982:1, 1983:3–1984:1, 1992:4–1995:3, and 1999:3–1999:4. In comparison to the Hamilton model, the 1978 period is no longer regarded as an upswing and the 1986 spike is clearly not an upswing. Instead, these periods are regarded as periods of high volatility around a low mean. The 1999–2000 and 1992–1996 upswings are also shorter than those determined by the Hamilton model.

The top panel of Figure 5 plots the squared deviations of the GDP growth rates from both the 11–quarter moving average trend and the HMM trend which is based on the entire dataset. Both methods clearly identify the mid 1990s as the lowest volatility period since 1978. Buckle, Haugh and Thomson (2001) suggest that the low volatility during this period was driven particularly by a temporary fall in the covariance across the sectors, which appears to cycle with

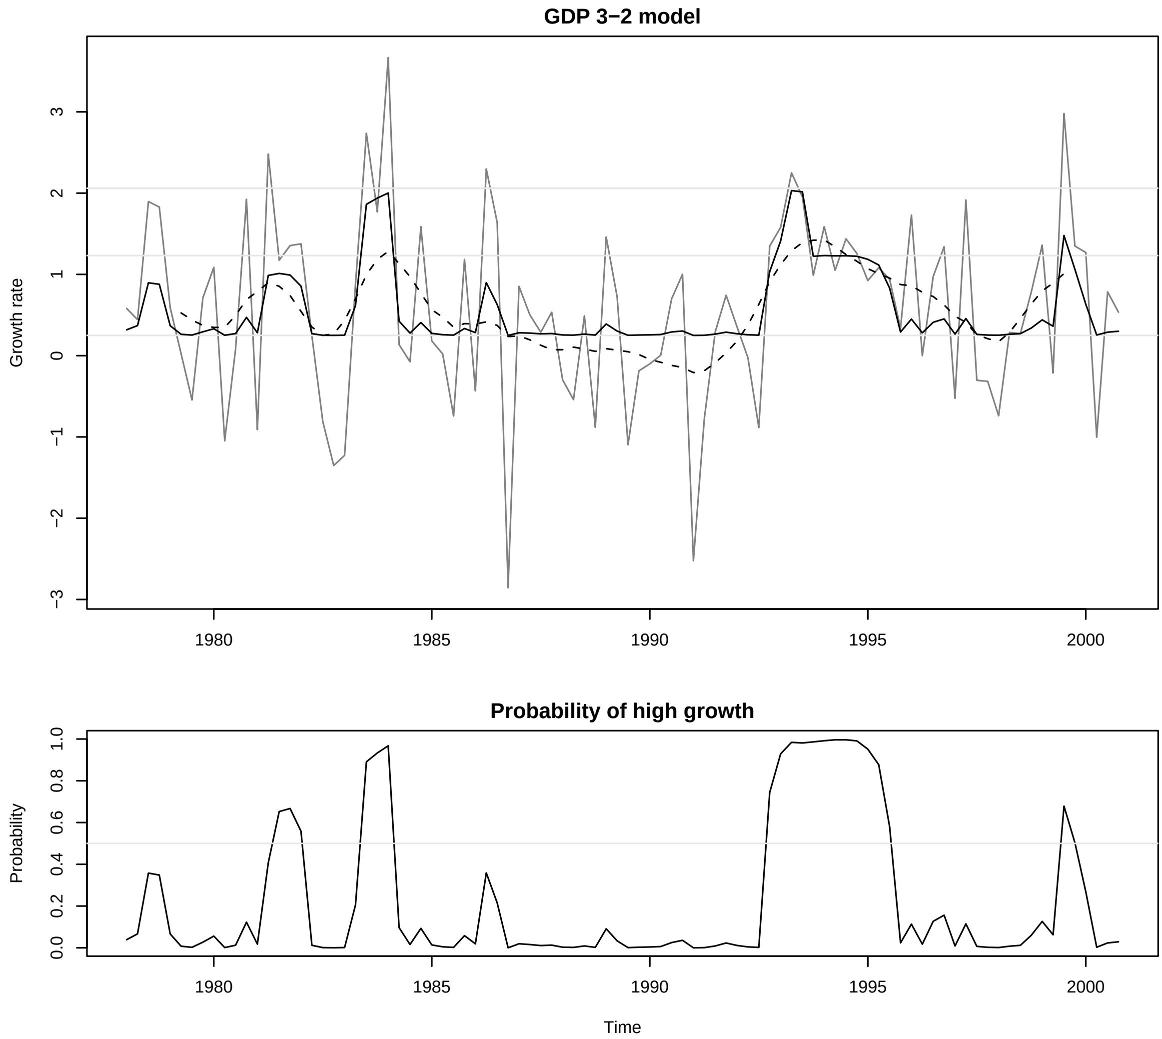

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	1.23	0.22	High growth	Low volatility
2	0.25	1.01	Low growth	High volatility
3	2.06	0.22	High growth	Low volatility
4	2.06	1.01	High growth	High volatility

Figure 4: Results of fitting a 3–2 model to quarterly GDP growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

no apparent trend. Interestingly, both periods of low volatility of New Zealand real GDP are periods when the economy switched to the high growth regime. The 1992:4 to 1995:3 period

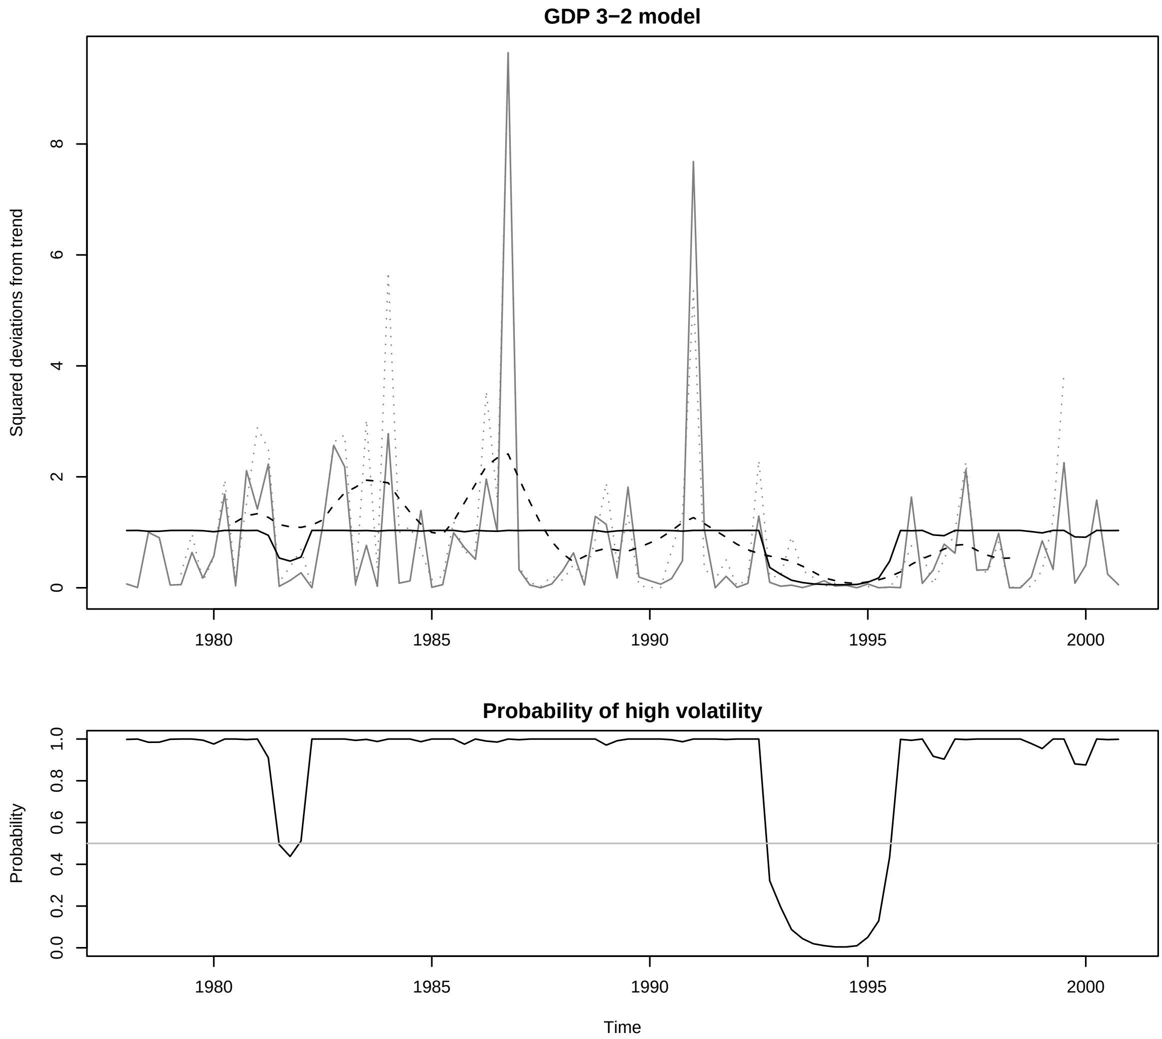

Figure 5: Results of fitting a 3–2 model to quarterly GDP growth rates. The top panel plots the squared deviations (grey dotted line) of the growth rates from their 11–quarter triangular moving average trend, and the squared deviations (solid grey line) of the growth rates from the HMM trend. The estimated volatility (black solid line) obtained from (15) and the triangular 11–quarter moving sample variance (black dashed line) are also plotted. The second panel plots the probability of being in the high volatility regime with the grey horizontal reference line equal to 0.5.

stands out however as a distinct period of nirvana, a period of high growth with low volatility.

The second panel of Figure 5 plots the probability of being in the high volatility regime and shows two periods during which the standard deviation switches from high to low volatility regimes. The estimated 3–2 model classifies most of the period between 1978 and 2000 as high volatility, with the possible exception of a short period from 1981:3–1982:1 and almost certainly a longer period from 1992:4 to 1995:3. As has been found for the United States (see Kim and Nelson, 1999; McConnell Perez–Quiros, 2000; Shaghil, Levin and Wilson, 2001) and several other developed economies including Australia (see Blanchard and Simon, 2001; and Simon, 2001), there is clear evidence of a switch to lower volatility of New Zealand real GDP during the 1990s. However, this switch to a lower volatility regime occurs much later than occurred in the US and Australia and, in contrast to the experience in these countries, the decline in volatility has not been sustained in New Zealand.

Plots of the probabilities of being in each of the four states given the data ($\gamma_t(j) = P(S_t = j|\mathbf{Y})$ for $j = 1, \ldots, 4$) are given in Figure 19 in the Appendix.

4.2 Services sector

The Services sector, as defined in Table 3, is the largest production sector and comprises approximately 50 percent of GDP. The MPQ version of the HMM model appears to be an appropriate characterisation of growth and volatility regimes experienced in this sector. The classification of states to regimes is shown in the table at the bottom of Figure 6. The two mean growth rates for the high growth regime are estimated at 1.25 percent per quarter in the low volatility regime and 1.67 percent per quarter in the high volatility regime. The two mean growth rates for the low growth regime are estimated at 0.68 percent per quarter in the low volatility regime and 0.03 percent per quarter in the high volatility regime.

The top panel of Figure 6 shows a plot of quarterly Services real output growth with the trend estimated from (14) and also from an 11–quarter triangular moving average. The probability of being in the high growth regime is plotted in the middle panel of Figure 6. The Services sector has experienced six periods between 1978 and 2000 when it switched from the low to the high growth regime. The periods in the high growth regimes are 1978:2–1978:4, 1981:2–1981:3, 1983:2–1984:2, 1985:4–1986:3, 1993:1–1995:3, 1998:4–1999:3. By comparison to GDP, there is clearer evidence of an upswing around 1986 in the Services sector suggesting that the effect of the introduction of GST inducing pre–spending is more marked in this sector which contains the wholesale and retail trade .

The Services sector has a clear and sustained break to lower volatility in 1992:1, as shown in Figure 7. This is the strongest evidence from any sector indicating a significant sustained downwards shift in volatility in the New Zealand economy. Buckle, Haugh and Thomson (2001) attribute this decline in Services volatility to declining volatility in the Finance and Real Estate industry and the Wholesale Trade industry.

This sustained decline in volatility in the Services sector is evident from deviations from both the 11–quarter centred moving average trend and deviations from the trend estimated by the MPQ model. The timing differs however. The moving sample variance indicates that the decline in volatility occurred around the end of the 1980s whereas Figure 7 shows the MPQ based variance switched to the lower volatility regime in 1992:1.

Plots of the probabilities of being in each of the four states given the data ($\gamma_t(j) = P(S_t = j|\mathbf{Y})$ for $j = 1, \ldots, 4$) are given in Figure 20 in the Appendix.

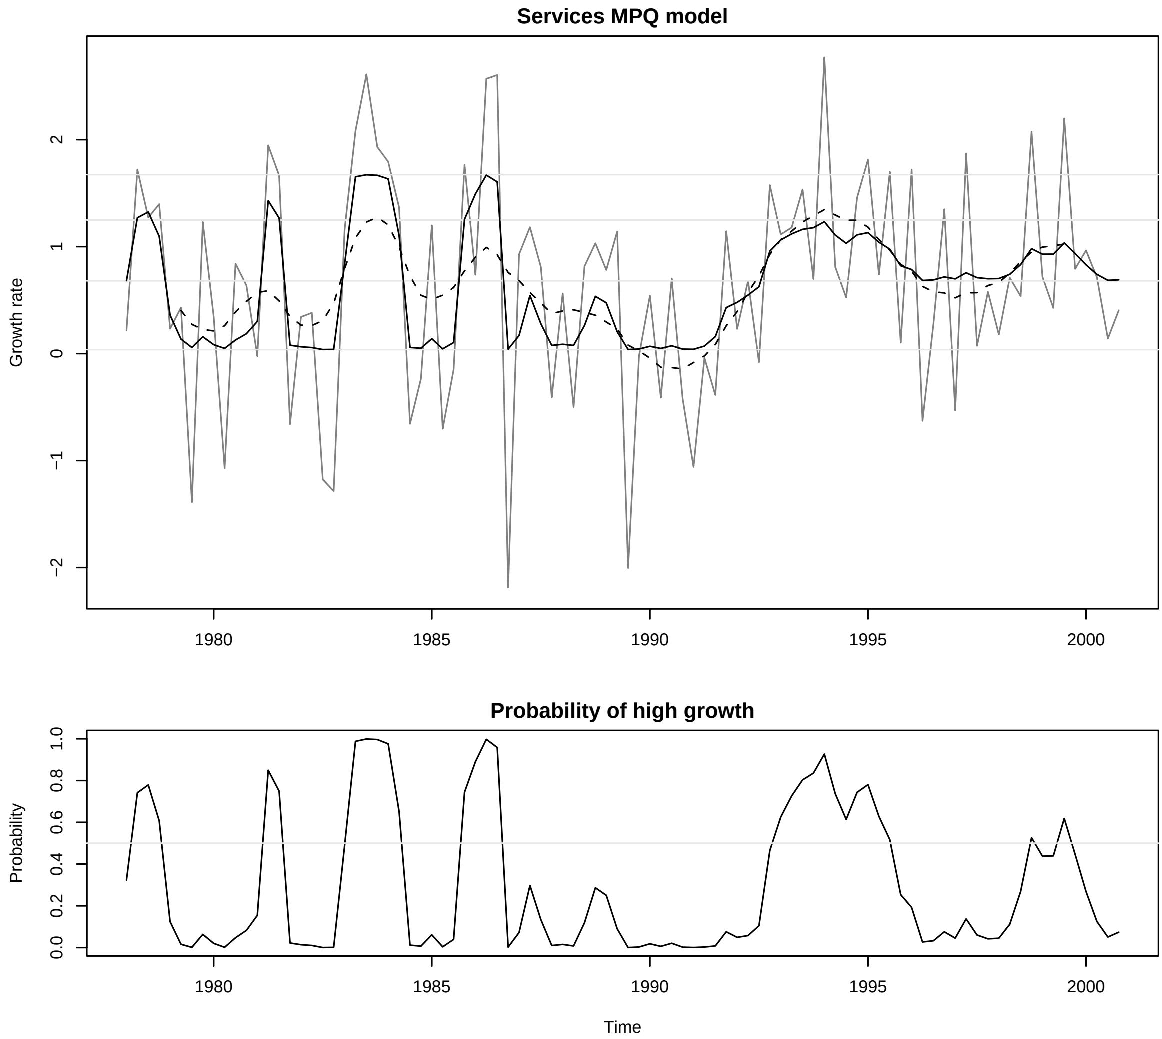

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	0.68	0.66	Low growth	Low volatility
2	0.04	0.86	Low growth	High volatility
3	1.25	0.66	High growth	Low volatility
4	1.67	0.86	High growth	High volatility

Figure 6: Results of fitting an MPQ model to quarterly Services growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

4.3 Government and Community Services sector

A two mean and two standard deviation model (2–2 model) was selected as an appropriate characterisation of the growth and volatility regimes experienced in the Government and Com-

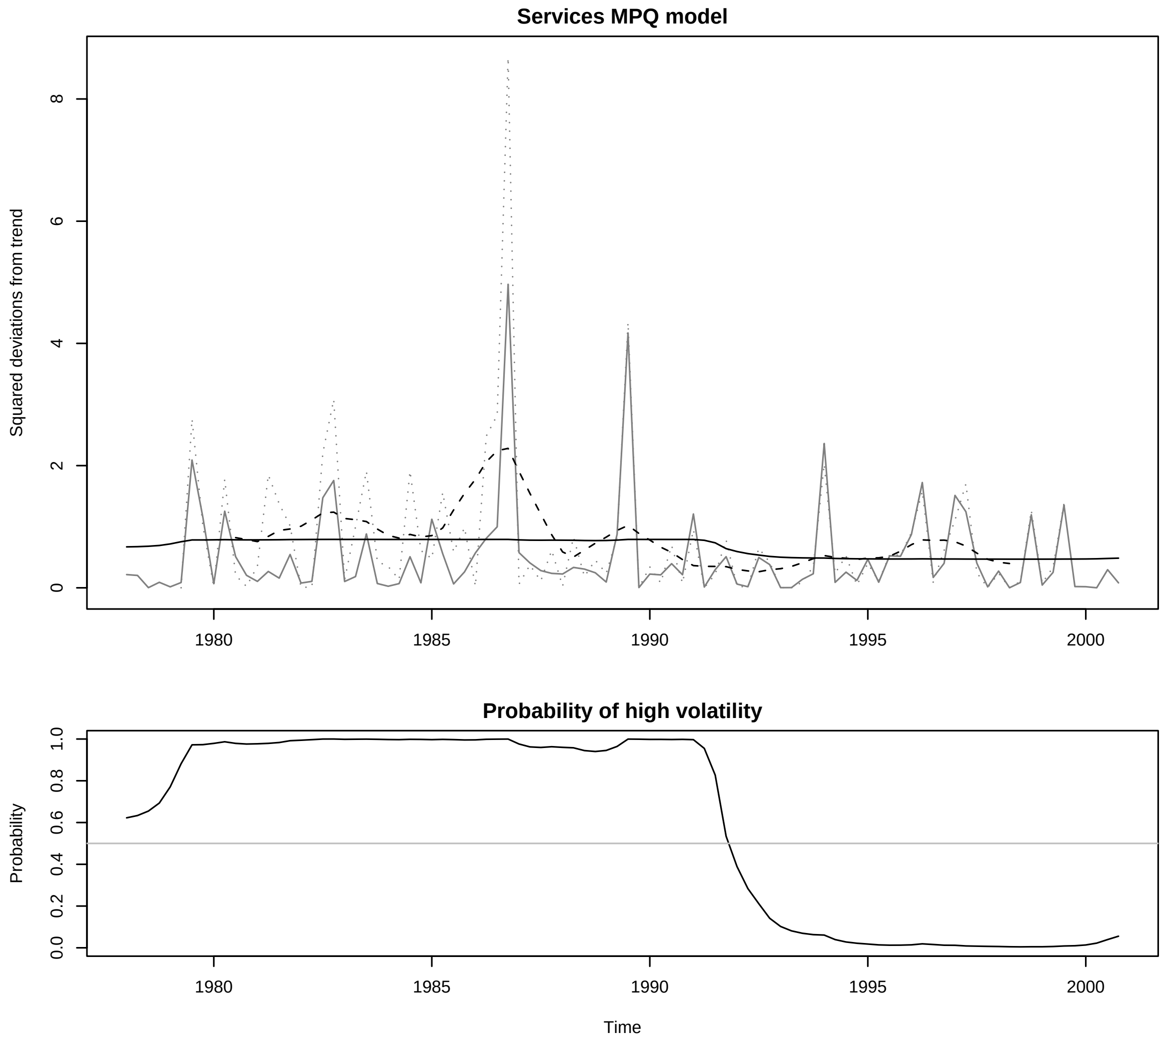

Figure 7: Results of fitting an MPQ model to quarterly Services growth rates. The top panel plots the squared deviations (grey dotted line) of the growth rates from their 11–quarter triangular moving average trend, and the squared deviations (solid grey line) of the GDP growth rates from the HMM trend. The estimated volatility (black solid line) obtained from (15) and the triangular 11–quarter moving sample variance (black dashed line) are also plotted. The second panel plots the probability of being in the high volatility regime with the grey horizontal reference line equal to 0.5.

munity Services sector. The classification of states to regimes is shown in the table at the bottom of Figure 8. In particular the high–growth mean rate is estimated to be 0.81 percent per quarter and the low–growth mean rate is estimated to be 0.28 percent per quarter.

This sector is characterised by relatively few growth and volatility regime switches and it has a tendency for sustained periods of time in one regime or another. Figure 8 illustrates that this sector switches from the high growth to the low growth regime in 1979:1 and remains in the low growth regime until 1992:3. At that time it switches to the high growth regime which it remains

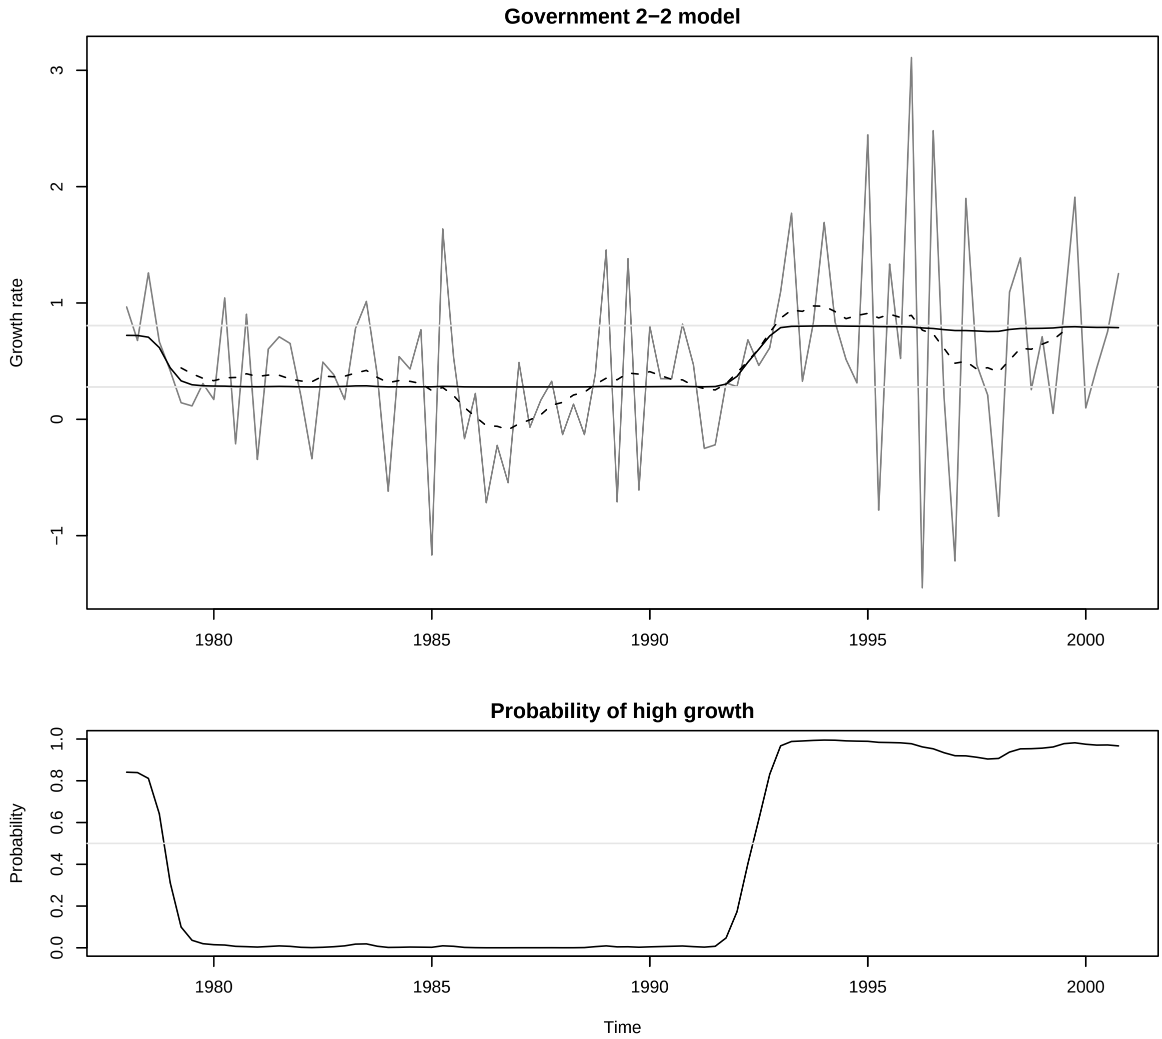

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	0.28	0.53	Low growth	Low volatility
2	0.28	1.08	Low growth	High volatility
3	0.81	0.53	High growth	Low volatility
4	0.81	1.08	High growth	High volatility

Figure 8: Results of fitting a 2–2 model to quarterly Government and Community Services growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

in for the remainder of the sample period. This switch to the high growth regime coincides with the transition of total GDP growth to the high growth regime at the beginning of the 1990s. In

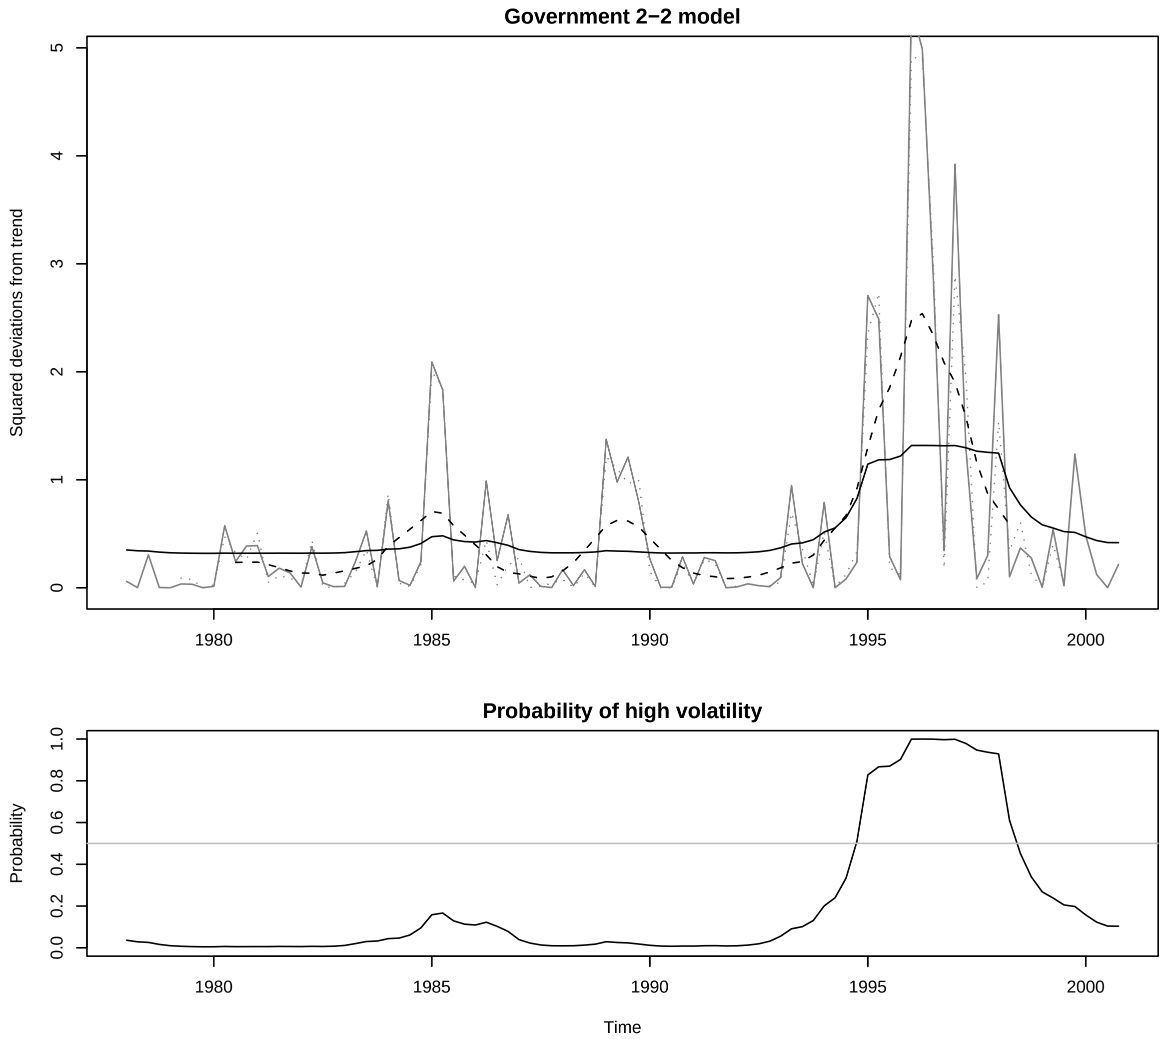

Figure 9: Results of fitting a 2-2 model to quarterly Government and Community Services growth rates. The top panel plots the squared deviations (grey dotted line) of the growth rates from their 11–quarter triangular moving average trend, and the squared deviations (solid grey line) of the GDP growth rates from the HMM trend. The estimated volatility (black solid line) obtained from (15) and the triangular 11–quarter moving sample variance (black dashed line) are also plotted. The second panel plots the probability of being in the high volatility regime with the grey horizontal reference line equal to 0.5.

contrast to total GDP growth, the Government and Community Services sector remains in this high growth regime.

This sector also experienced a switch to the high–volatility regime in 1994:1, as shown in Figure 9. The reasons for this regime switch are unknown. Plots of the probabilities of being in each of the four states given the data ($\gamma_t(j) = P(S_t = j|\mathbf{Y})$ for $j = 1, \ldots, 4$) are given in Figure 21 in the Appendix.

4.4 Manufacturing sector

The Manufacturing sector is difficult to characterise with only one version of the HMM model. This sector appears to be characterised by two means and a constant standard deviation for most of the sample period, but the early 1980s appear to have a different structure. There was a very strong surge in manufacturing output growth in 1983–1984 which can be regarded as either a high volatility period (which would require a model which allowed for a low and high volatility states, such as the MPQ model) or as a period with very high mean growth (so that a model with one volatility state, but at least three mean states would be required).

Three models were ultimately selected to analyse the Manufacturing sector: the Hamilton model (2–1 model), the MPQ model (4–2 model), and a three mean and constant standard deviation model (3–1 model). The Hamilton model seems the most appropriate for this sector with the exception of the early 1980s period.

Figure 10 shows the results of estimating the Hamilton model for the Manufacturing sector. The high–growth mean is estimated as 1.89 percent per quarter and the low–growth mean as -0.73 percent per quarter with a constant standard deviation of 1.77 percent per quarter. The Hamilton trend closely tracks the 11 quarter triangular moving average although the Hamilton model is more flexible and is able to identify more turning points. An 11 quarter uniform moving average was unable to find many of the turning points picked out by both the Hamilton trend and the triangular moving average. The Hamilton model picks out 6 periods when the Manufacturing sector switched from the low growth to the high growth regime: 1978:1–1979:2, 1980:4–1982:1, 1983:2–1984:4, 1991:3–1992:1 and 1992:4–1995:1, 1999:1–1999:4. There are also short switches to high growth in 1986–1987, 1989, 1996 and 1997, but the probabilities only just exceed 50 percent and last only one quarter.

The Hamilton model is parsimonious with only 6 parameters and is superior to both the MPQ and 3–1 models on the basis of both the AIC and BIC model selection criteria. However, it does not seem to be as appropriate for the early 1980s period. This period is different probably because of the “Think Big” capital expenditure program which boosted activity in the Manufacturing sector, notably the Machinery and Equipment Manufacturing industry (see Buckle, Haugh and Thomson (2001) who identify the Machinery and Equipment manufacturing industry as the key source of the strong surge in volatility in the early 1980s). Including the early 1980s period in the estimation of the Hamilton model is likely to result in an upwards bias for the estimated means for the subsequent periods. Estimates using the MPQ and the 3–1 models suggest that this is the case.

Using the MPQ model, Manufacturing can be characterised by the four states and corresponding regimes shown in the table at the bottom of Figure 14. Furthermore Figure 11 implies that Manufacturing is in the low volatility regime with a standard deviation of 1.23 percent per quarter for most of the sample period except for the period 1982:1– 1984:1 when it switches to the high volatility state with a standard deviation of 1.89 percent per quarter. For this period the overall volatility of the data is also higher since the mean growth rates of 3.83 percent and -3.07 percent per quarter produce the most extreme deviations about the constant overall mean. This is consistent with the triangular moving sample variance which is also at its highest around this period and at a clearly higher level than the rest of the sample.

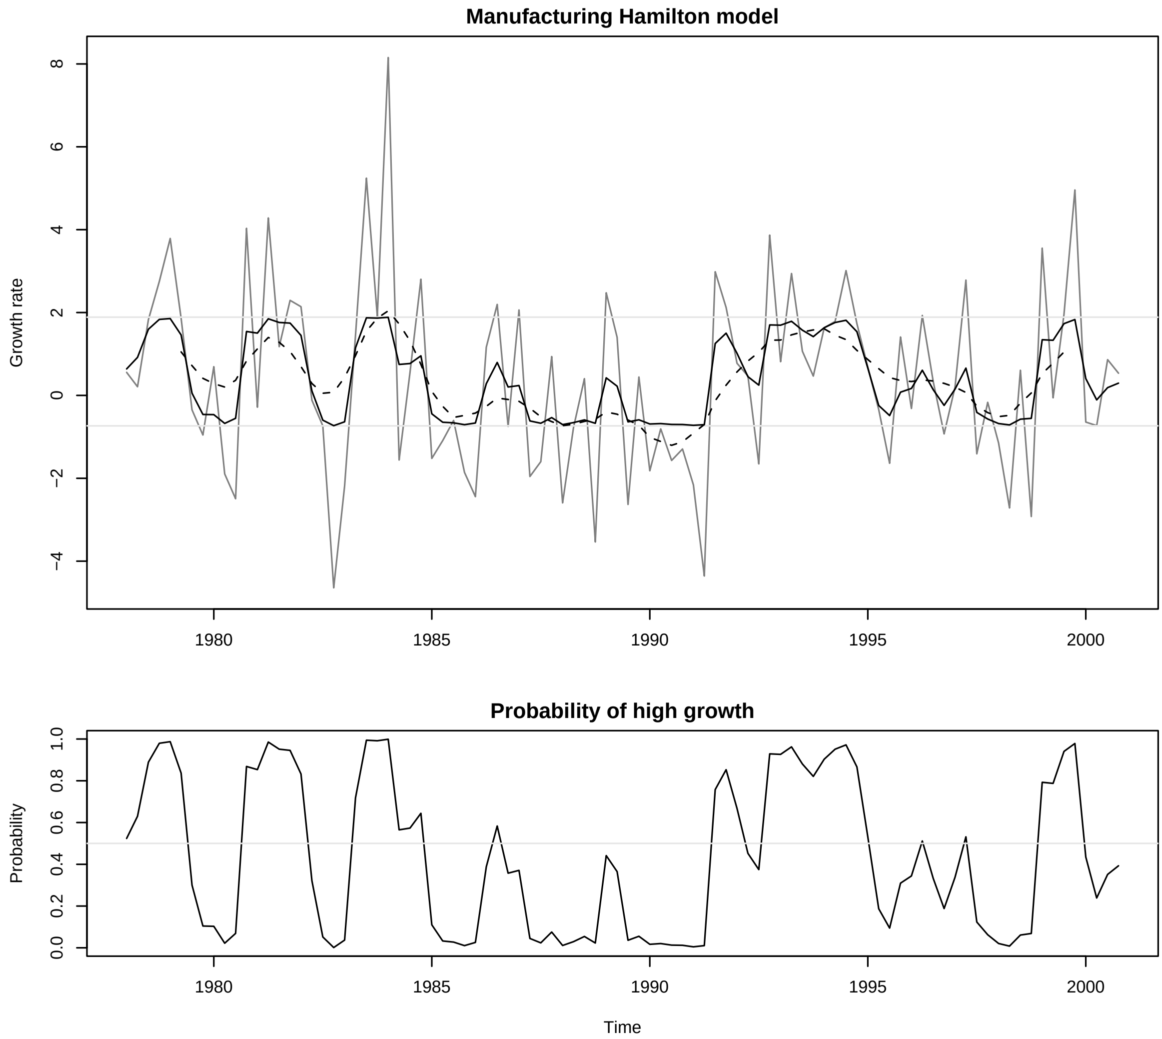

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1, 2	-0.73	1.77	Low growth	Constant volatility
3, 4	1.89	1.77	High growth	Constant volatility

Figure 10: Results of fitting a Hamilton model to quarterly Manufacturing growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

Outside this period of high volatility the high and low mean growth rates of 1.74 percent and -1.01 percent per quarter are closer to those for the Hamilton model. The MPQ values are however slightly lower which leads to more regime switching because there is now a lower threshold before growth switches to the high regime. As a result the MPQ model suggests that

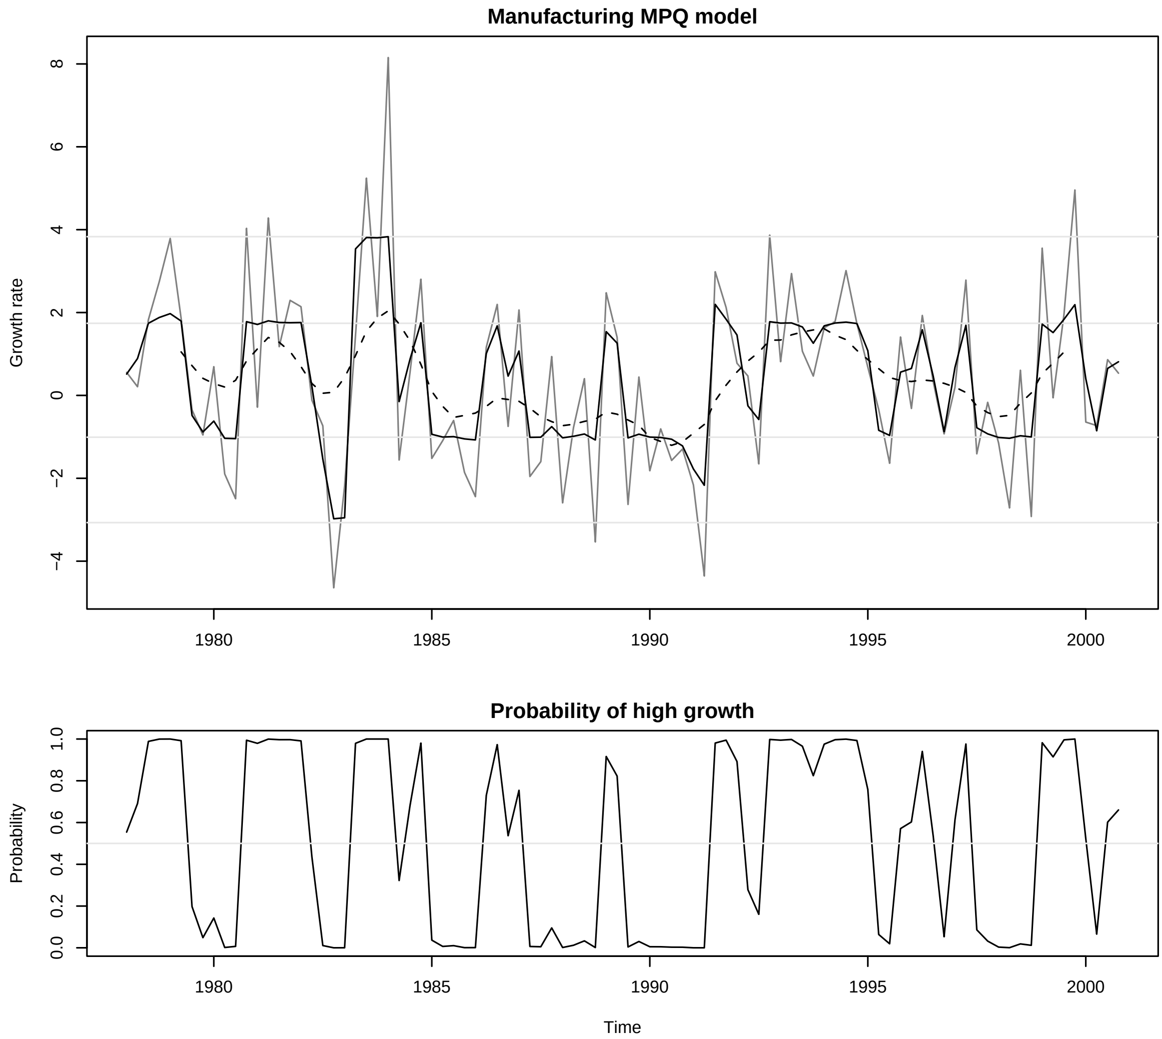

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	-1.01	1.23	Low growth	Low volatility
2	-3.07	1.89	Low growth	High volatility
3	1.74	1.23	High growth	Low volatility
4	3.83	1.89	High growth	High volatility

Figure 11: Results of fitting an MPQ model to quarterly Manufacturing growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

there are nine rather than five switches from low to high growth regimes, some of which even the 11–quarter triangular moving average is unable to track. These regimes last for periods of

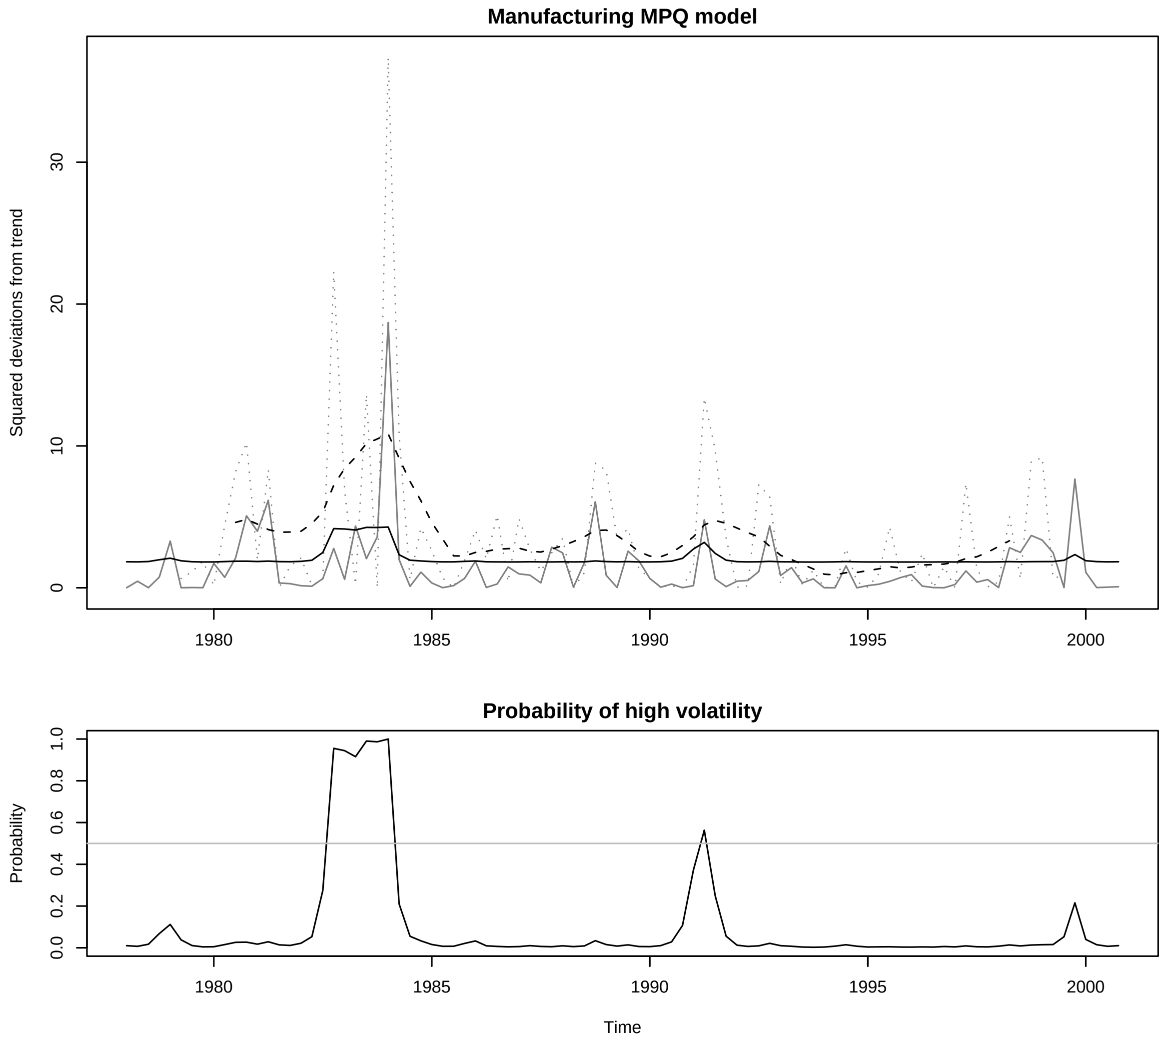

Figure 12: Results of fitting an MPQ model to quarterly Manufacturing growth rates. The top panel plots the squared deviations (grey dotted line) of the growth rates from their 11–quarter triangular moving average trend, and the squared deviations (solid grey line) of the GDP growth rates from the HMM trend. The estimated volatility (black solid line) obtained from (15) and the triangular 11–quarter moving sample variance (black dashed line) are also plotted. The second panel plots the probability of being in the high volatility regime with the grey horizontal reference line equal to 0.5.

between two and ten quarters which makes it difficult for any one moving average to detect all of them. The timing of the upswings identified by the MPQ model are generally very close to those identified by the Hamilton model. However the MPQ model more clearly identifies the upswings to high growth regimes in 1986–1987, 1996 and 1997 and clearly identifies another upswing beginning in 2000:3.

The 3–1 model (illustrated in Figures 12 and 13) also isolates the early 1980s period as different from the rest of the sample, but uses a third mean level to do this. The three means estimated

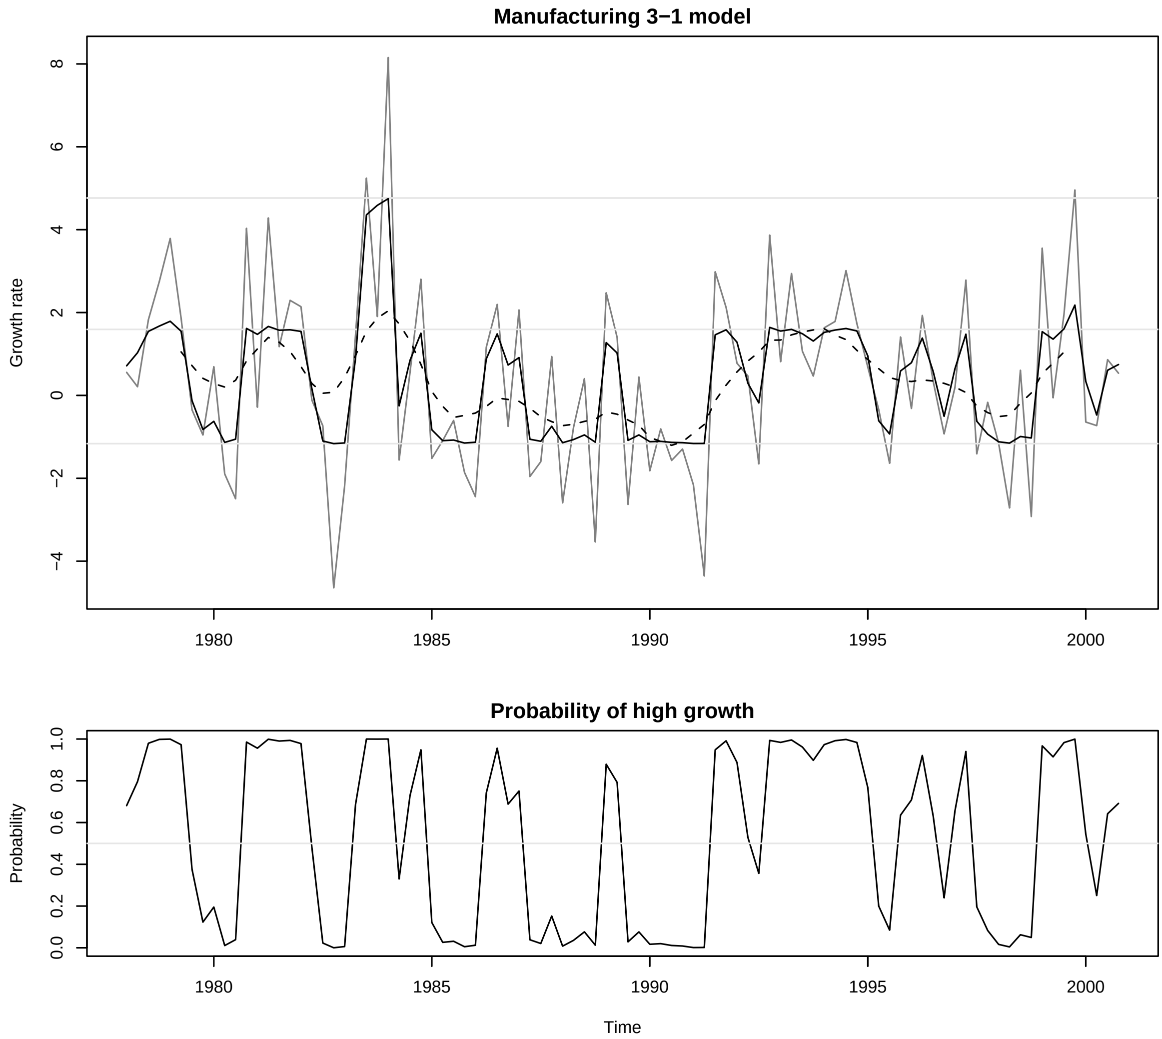

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	-1.16	1.45	Low growth	Constant volatility
2	1.59	1.45	High growth	Constant volatility
3, 4	4.76	1.45	High growth	Constant volatility

Figure 13: Results of fitting a 3–1 model to quarterly Manufacturing growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

by this model are -1.16 percent, 1.59 percent and 4.78 percent per quarter, all with a standard deviation of 1.45 percent per quarter. The high growth state with a mean of 4.78 percent growth is realised once, from 1983:3 to 1984:1. During the rest of the sample period Manufacturing switches between the low mean of -1.16 percent and the high mean of 1.59 percent per quarter, which are also slightly lower than the two Hamilton means. The timing of switches to the 1.59 percent high growth state is very similar to the timing of switches derived using the MPQ model and, as with the MPQ model, the 3–1 model has more switches to the high growth regime than the Hamilton model. The 3–1 model (where nine parameters are estimated) is superior to the MPQ model on the basis of the AIC and BIC criteria because it has fewer parameters to estimate than MPQ (eleven parameters).

Overall, the manufacturing sector has had frequent regime switches with varying lengths. Interpretation of the number of switches depends on how the early 1980s period is treated. There appears to be no permanent change in volatility in this sector, with the possible exception of a brief spike in volatility in the mid 1980s.

Again plots of the probabilities of being in each of the four states given the data ($\gamma_t(j) = P(S_t = j|\mathbf{Y})$ for $j = 1, \ldots, 4$) are given in Figures 22 and 23 in the Appendix for both the MPQ and 3–1 models.

4.5 Primary sector

A model with three mean growth rates and two standard deviations (a 3–2 model) seems to provide a good characterisation of the Primary sector. Further details are given in Figures 14, 15 and 24. Although the MPQ model also fits the data well, two of the estimated four means were very close to each other (0.80 percent and 1 percent) and, with one less parameter to estimate, a 3–2 model with three states had the better BIC value.

Figure 15 shows that, throughout the sample period, the Primary sector was predominantly in the high volatility regime with a 0.97 percent mean quarterly growth rate. This state prevailed from 1978:1 to 1988:1 and from 1994:1 to 2000:4. This tendency to remain in one regime for very long periods distinguishes the Primary sector from other sectors and aggregate GDP, with the exception of the Government and Community Services sector. The Primary sector also illustrates how an informal analysis of the raw quarterly growth rate data could be misleading. The high volatility in the data may make it tempting to infer from a new large observation that there has been a shift in the mean and a transition to a new part of the cycle. The 3–2 model suggests this would often be an incorrect interpretation for this sector.

During the middle of the sample period from 1988:2–1994:1 the Primary sector switches to the low volatility regime and exhibits two cycles that switch between the lowest (-1.74 percent per quarter) and highest (4.88 per cent per quarter) mean growth states. During this period shifts in the mean growth rates, rather than changes in the standard deviations, drive changes in the overall volatility of Primary output.

In summary, the Primary sector temporarily switches to low volatility in the late 1980s and early 1990s and returns to the high volatility regime with a constant mean from 1994 onwards. This implies that the Primary sector has not made a permanent contribution to the decline in GDP volatility, a conclusion that is also made in Buckle, Haugh and Thomson (2001).

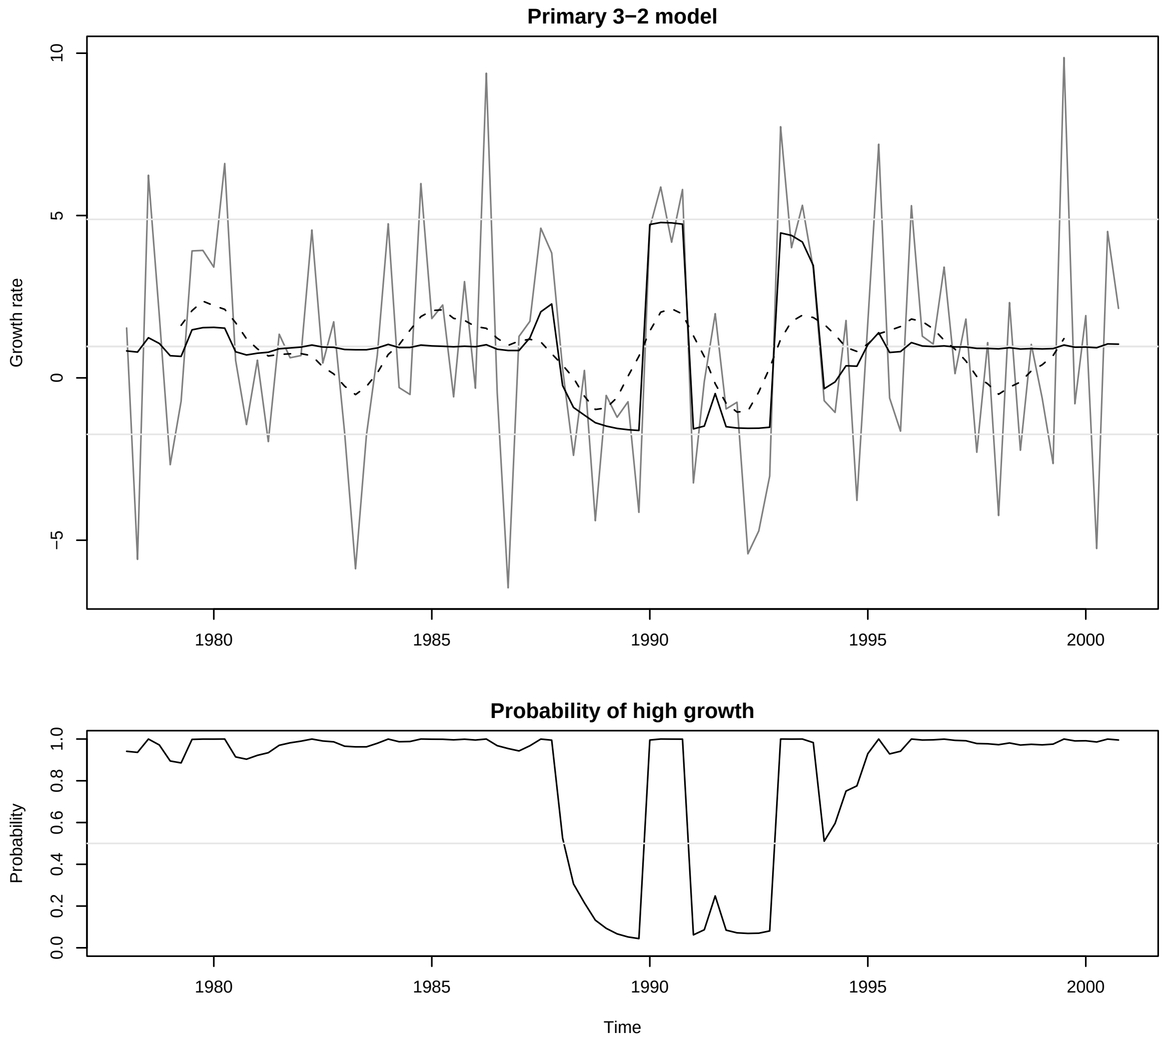

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	-1.74	1.89	Low growth	Low volatility
2	4.88	1.89	High growth	Low volatility
3, 4	0.97	3.30	High growth	High volatility

Figure 14: Results of fitting a 3–2 model to quarterly Primary growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

4.6 Construction sector

The Construction sector is the most volatile of New Zealand's five production sectors and was a difficult sector to model and characterise. A 3–1 model was preferred over other possible

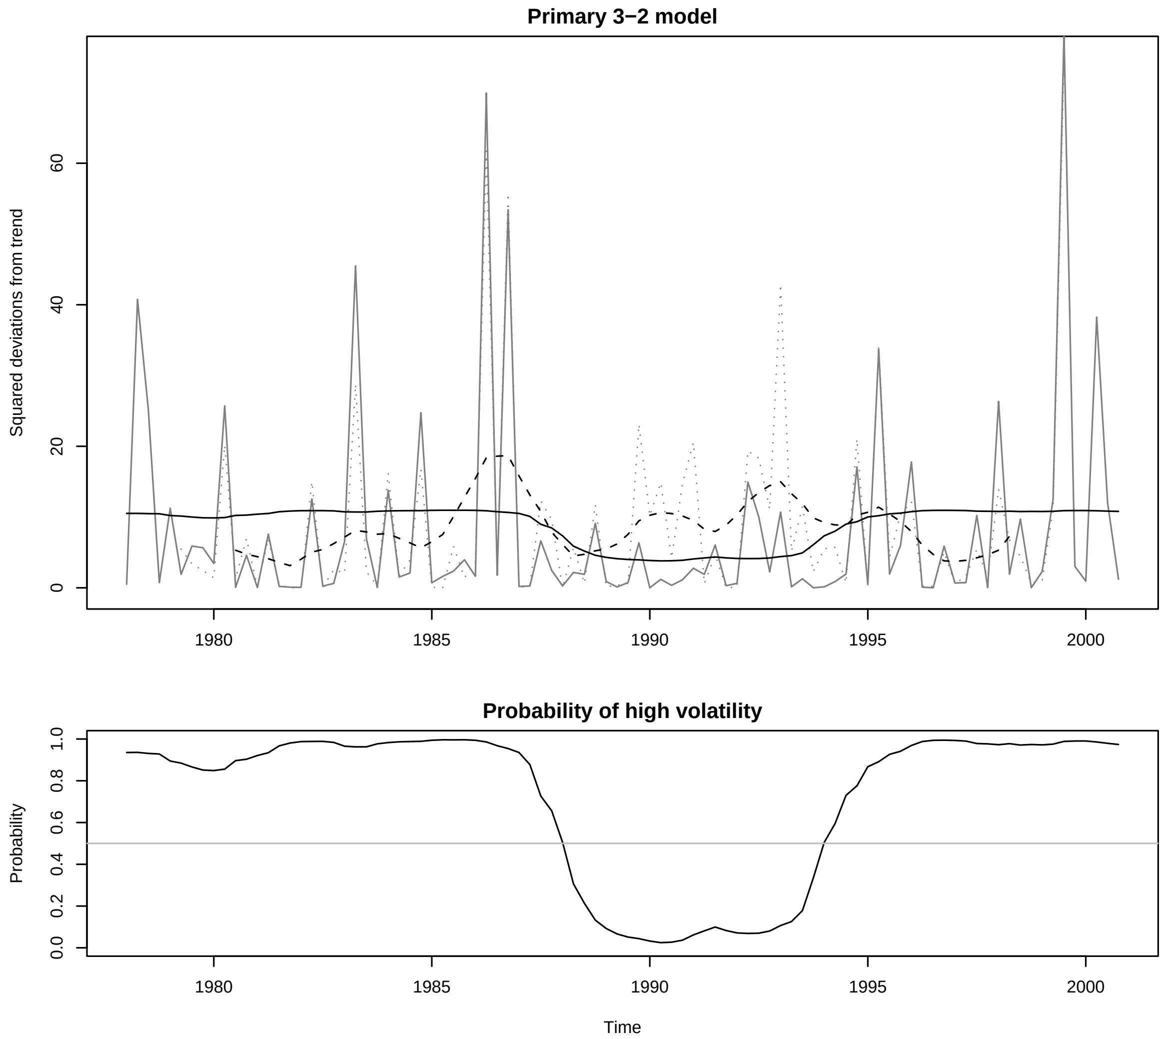

Figure 15: Results of fitting a 3–2 model to quarterly Primary growth rates. The top panel plots the squared deviations (grey dotted line) of the growth rates from their 11–quarter triangular moving average trend, and the squared deviations (solid grey line) of the GDP growth rates from the HMM trend. The estimated volatility (black solid line) obtained from (15) and the triangular 11–quarter moving sample variance (black dashed line) are also plotted. The second panel plots the probability of being in the high volatility regime with the grey horizontal reference line equal to 0.5.

model structures. The three mean quarterly growth rates are -10.29 percent, -1.3 percent and 2.66 percent and the standard deviation is 3.98 percent per quarter. The state with -10.29 percent mean growth rate (state $S_t = 1$) is, in effect, an outler state since it picks off extreme low growth rates. (Other outlier models might also be considered for this data since there also appear to be a few extreme high growth rates.) For most of the sample period the Construction sector switches between a low–growth mean of -1.3 percent mean and a high–growth mean of 2.66 percent.

Figure 16 implies that the Construction sector switches frequently between high and low growth

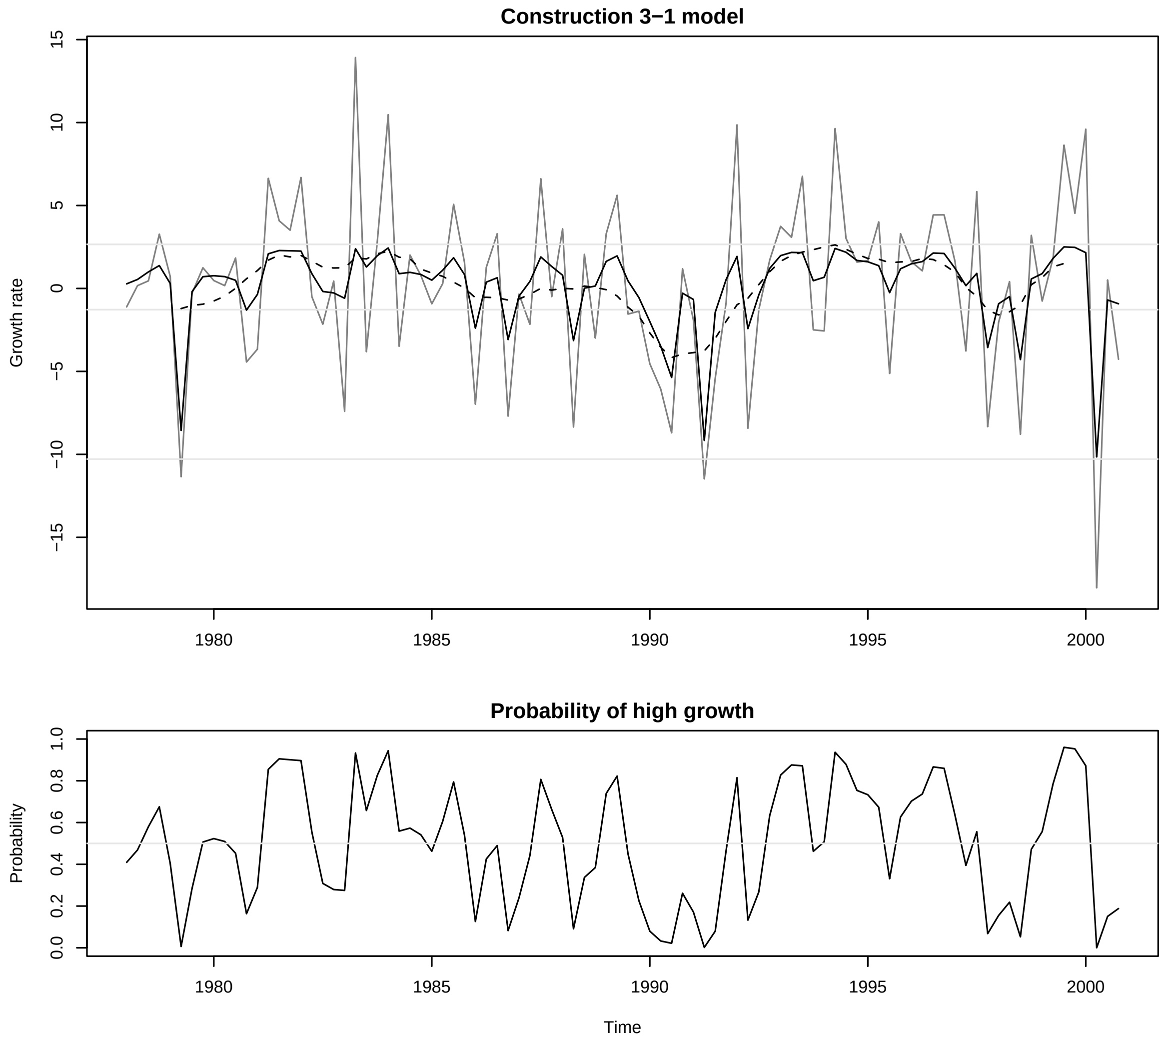

S_t	μ_{S_t}	σ_{S_t}	Regime classification	
1	-10.29	3.98	Outlier low growth	Constant volatility
2	-1.30	3.98	Low growth	Constant volatility
3, 4	2.66	3.98	High growth	Constant volatility

Figure 16: Results of fitting a 3–1 model to quarterly Construction growth rates. The top panel shows the growth rates (grey line) with the trend (solid line) estimated from (14) and also from an 11–quarter triangular moving average for comparison (dashed line). The grey horizontal lines represent the estimated μ_j. The second panel plots the probability of being in the high growth regime with the grey horizontal reference line equal to 0.5. Estimated mean growth rates and standard deviations for each state, and the classification of states to regimes are shown in the bottom panel.

regimes. If the quarters 1985:1, 1993:4 and 1995:3 are disregarded as switches to the low growth state (the probabilities of these quarters being in the low growth state are only 0.54, 0.54 and 0.67 respectively) there are 10 upswings from the low to the high growth regime during the 23 years of the sample (there are nine if the one quarter duration upswing in 1992:1 is ignored).

5 Dating New Zealand business cycles

One of the striking conclusions that emerged from Hamilton's early research was the close relationship between the NBER dating of US business cycle booms and recessions and the timing of switches between low and high growth regimes identified by the Markov switching model. This suggested the Hidden Markov modelling procedure could be used as an alternative procedure for dating business cycles. An example of this type of application is the work of Krolzig (1997) who has used the HMM technique to date turning points in the business cycles of Australia, Canada, France, Germany, Japan, United Kingdom and the United States and to evaluate the relationships between the business cycles in these countries. This section compares the dates the New Zealand economy switched between low and high growth regimes as identified by the HMM models fitted to New Zealand real GDP data (using the Hamilton model and the 3–2 model) with previous estimates of New Zealand business cycle dating points obtained by more traditional methods.

Business cycles are typically measured as "classical" cycles or as "growth" cycles. The measurement of the classical cycle is concerned with the identification of expansions and contractions in absolute levels of economic activity. Its origins are in the work of Mitchell (1913, 1927) and Burns and Mitchell (1946). This approach defines the cycle as the co–movement of a large number of economic activities with many activities expanding and then contracting together. Its tradition has been continued as the "NBER approach" and formalised by an algorithm developed by Bry and Boschan (1971).

Growth cycles on the other hand are movements in the detrended level of economic activity. This approach uses filters such as the Henderson Moving Averages (HMA), the Hodrick–Prescott (HP) filter or equivalent formulations based on structural time series models (see Harvey (1989) and Harvey and Jaeger (1993)) to remove a trend from the data. The resulting trend deviations are then analysed to date the growth cycle using a decision rule such as 2 quarters of negative growth followed by a quarter of positive growth defining a trough for example.

Kim, Buckle and Hall (1995) have comprehensively reviewed early applications of the NBER approach to New Zealand GDP data. They also apply the Bry and Boschan algorithm and compared the business cycle dates generated by that method with growth cycle dates they generated by applying the Henderson Moving Averages (HMA), Hodrick–Prescott (HP) and structural time series (Harvey (1989)) trend estimation procedures. The methods of Harvey (1989) are implemented in the computer package STAMP (Koopman et al (1995)). Table 4 summarises the business cycle turning points estimated by Kim, Buckle and Hall (1995) and compares them with the dates for high and low growth regimes generated by the hidden Markov model for total GDP growth rates (using the Hamilton model and the 3–2 model). To make this comparison, a period characterising a movement from a trough to a peak of the business cycle is regarded as the period during which a high growth regime prevailed, and movement from a peak to a trough is regarded as the period during which a low growth regime prevailed.

The NBER and HMM methods are attempting to identify movements in the actual and trend levels of economic activity and therefore should be broadly comparable, except where an actual observation has a large noise component which the HMM method will disregard when determining the cycle, but which the classical definition will not. The deviations from trend methods on the other hand regard deviations from the trend level as the basis for identifying turning points in the cycle. Nevertheless, the deviations from trend method will identify similar turning points to HMM if a smooth long–run trend is fitted to the levels series. Once a long–run trend is fitted and a deviations from a trend series is estimated, peaks and troughs in this deviations series can be used to identify points at which there is a switch from one growth state to another. For example, assume an economy with 2 growth states, high and low. If the economy is in the high–growth state there will be a positive deviation from the long–run levels trend, which will be growing at a rate somewhere in between the two local high–growth and low–growth states. At the end of the high–growth state the deviation from the long–run level will be at its peak. Once the economy switches to the low growth state the level will start to decline towards the long–run levels trend again. This is illustrated in Figure 1.

Turning Points	**HMM Ham**	**HMM 3–2**	**STAMP & HP**	**HMA**	**BB & NBER**
Trough	1978:1		1978:1	1978:1	1978:1
Peak	1978:4		1980:1	1979:1	
Trough	1981:1	1981:2	1981:1	1980:4	
Peak	1982:1	1982:1	1982:2	1982:1	1982:2
Trough	1983:2	1983:2	1983:1	1983:1	1983:1
Peak	1984:2	1984:1	1984:1	1984:2	
Trough	1986:1		1986:1	1985:4	
Peak	1986:2		1986:3	1986:3	1986:3
Trough			1988:2	1988:3	
Peak			1989:2	1990:3	
Trough	1992:3	1992:3	1991:2	1991:3	1991:2
Peak	1996:1	1995:3			
Trough	1998:4	1999:2			
Peak	2000:2	1999:4			

Table 5: Turning points for real GDP growth as determined by hidden Markov models, deviation from trend and NBER methods. The HMM models used were the Hamilton model (Ham) and the 3–2 model. The turning points from STAMP, the Hodrick–Prescott filter, the Henderson moving Average (HMA) and the Bry and Boschan methods (BB) are taken from Kim, Buckle and Hall (1995), Table 4, page 162. The dates derived using STAMP and the Hodrick–Prescott filter were exactly the same apart from a one quarter difference in the timing of the 1978 trough. Hence, for brevity, this table lists the STAMP and Hodrick–Prescott dates in the same column and uses the STAMP result to date the 1978 trough. Kim, Buckle and Hall's results are based on Statistics New Zealand's seasonally–adjusted fixed–weight production GDP series for the period 1977:2 to 1993:2. The HMM (Hamilton and 3–2 models) results are derived for the calibrated seasonally–adjusted chain–linked production GDP data covering the period 1977:2 to 2000:4.

It is clear from Table 4 that the Hamilton version of the hidden Markov model produces turning point dates that are very similar to those identified by Kim, Buckle and Hall (1995) using alternative deviation from trend methods. The dates for the Hamilton model based are less similar to those generated by the Bry and Boschan algorithm that represents the NBER method. However the Bry and Boschan algorithm appears to miss a distinct peak in 1984 and therefore regards the 1986 GST spike as the peak for the high growth period commencing in 1983:1.

One noticeable difference between the dates for the Hamilton model and those for the deviation from trend approach is the identification of an early 1990s trough. The deviation from trend methods date it as 1991:2 (or 1991:3) whereas the Hamilton dates the trough at 1992:3. This result supports the idea that the hidden Markov model is better at distinguishing a change in growth than traditional methods. In the first half of 1991 there were two negative quarters followed by two positive quarters and then a return to two negative quarters. The traditional methods date the trough as the end of the first two negative quarters while the Hamilton model dates the turning point as after the second two negative quarters.

The 3–2 model does not regard the mid 1986 GST spike as a peak but rather volatility around a low trend. The model has two standard deviations that give it a flexibility to interpret an observation as high growth in a low volatility regime or low growth in a high volatility regime. The deviation from trend method and Hamilton model only determine whether a new observation is the continuation of the current mean growth rate or the start of a new mean growth rate as opposed to determining whether there has been a switch in volatility.

The hidden Markov model can draw finer distinctions between different parts of a series than either the Bry and Boschan procedure or the deviation from trend methods. This is an advantage in dating cycles because it has more flexibility to determine when there has been a distinct turning point. An observation may be different from previous ones but, rather than being a turning point in the cycle, it may only signal a change in volatility (3–2 model) or it may not be large enough to constitute a change of regime (Hamilton). For more standard points the HMM and deviations from trend methods have very similar turning points despite their different conceptual bases.

6 Growth and volatility regimes across sectors

The estimation of Markov switching models for production sectors provides a basis for interpreting the temporal relationship between growth and volatility regimes across the sectors and their contribution to the timing of regime switching evident in total GDP. Although this issue could be more rigorously pursued using a vector Markov switching model of the type suggested by Kontomelis (2001) for example, some useful insights can nevertheless be gained by an informal analysis of the relationship between growth and volatility regimes across sectors and their relationship with total GDP growth and volatility.

Figure 17 plots and compares the probability of each production sector and total GDP being in a high growth regime for each quarter from 1978:1 to 2000:4. The vertical lines on Figure 17 represent the dates that total GDP switches from a low to a high growth regime and vice versa. Table 5 dates the periods when total GDP, Services, Manufacturing, and Construction were in

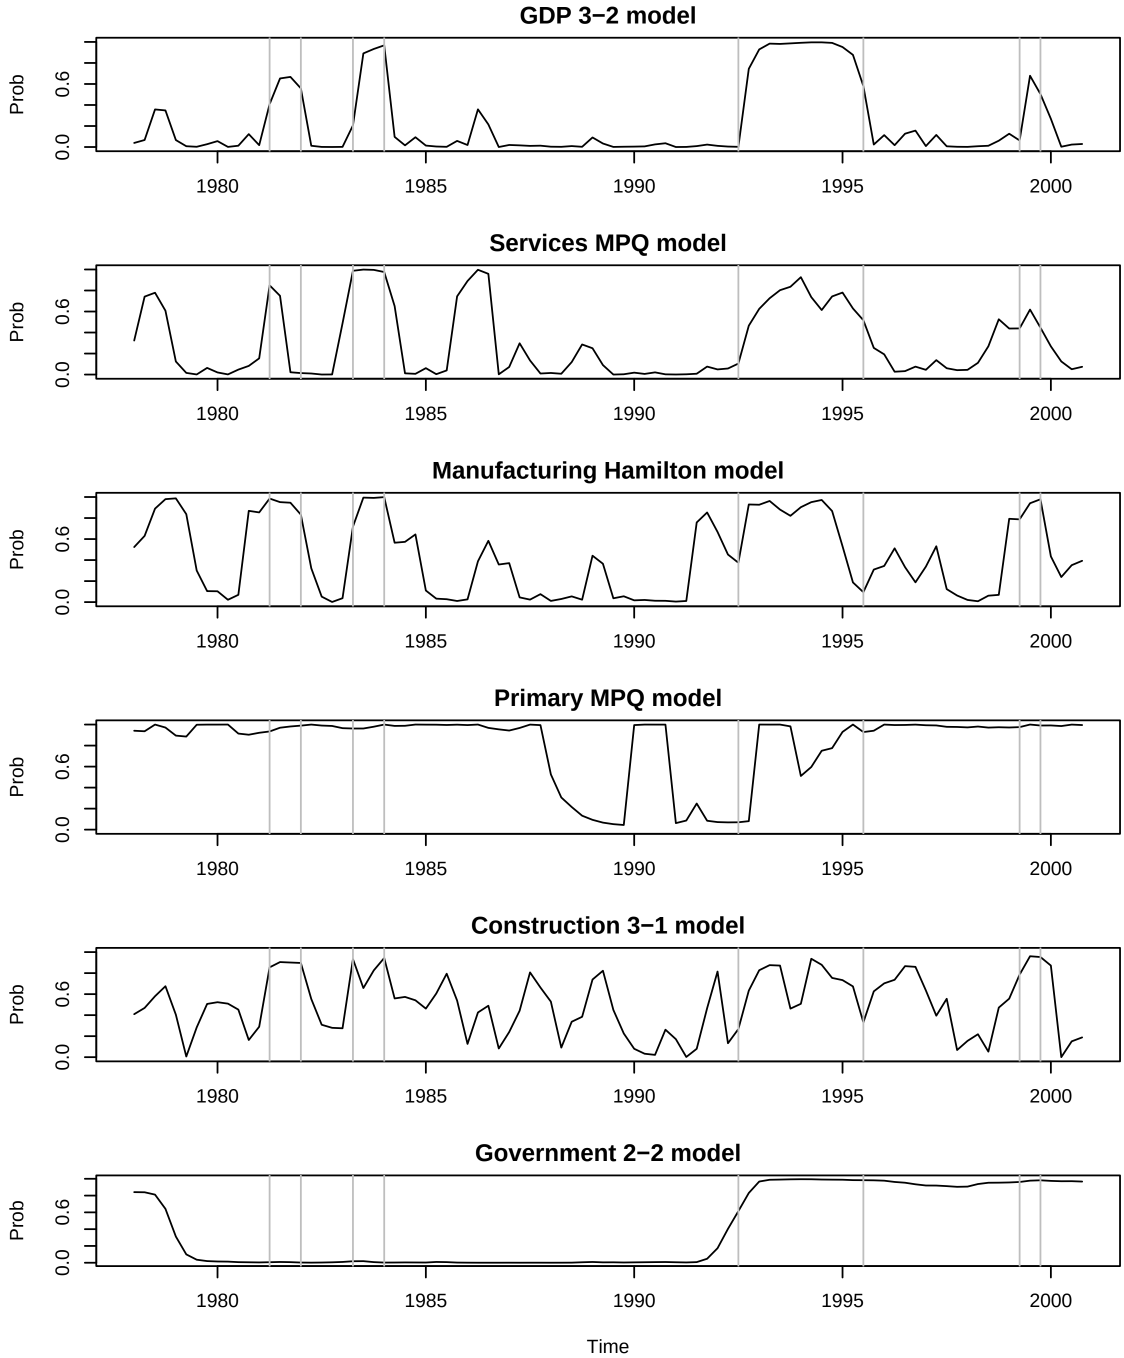

Figure 17: Probability of each production sector and total GDP growth rates being in a high growth regime for each quarter from 1978:1 to 2000:4. The vertical grey lines represent the dates that total GDP switches from a low to a high growth regime and vice versa.

the high growth regimes. The Primary sector and Government and Community Services sector are not included in Table 5 because they display far fewer growth regime switches. The Primary sector is in the high growth regime during 1978:1 to 1988:1, 1990:1 to 1990:4 and 1993:1 to

2000:4 (that is, the only periods of low growth in the Primary sector are 1988:2–1989:4 and 1991:1–1992:4). The Government and Community Services sector only shows one switch from low to high growth in 1992:3 and remained continuously in the high growth regime from 1992:3 to 2000:4.

Data	Model	Dates of high growth regimes				
GDP	Ham	1978:2–1978:4	1981:2–1982:1	1983:3–1984:2	1992:4–1996:1	1999:1–2000:2
Ser	MPQ	1978:2–1978:4	1981:2–1981:3	1983:2–1984:2 1985:4–1986:3	1993:1–1995:3	1998:4–1999:3
Man	Ham	1978:1–1979:2	1980:4–1982:1	1983:2–1984:4	1991:3–1992:1 1992:4–1995:1	1999:1–1999:4
Con	3–1	1978:3–1978:4 1979:4–1980:2	1981:2–1982:2	1983:2–1985:4 1987:3–1988:1 1989:1–1989:2	1992:4–1997:1	1999:1– 2000:1

Table 6: Dates when total GDP, Services (Ser), Manufacturing (Man), and Construction (Con) were in the high growth regime. The models fitted are as indicated with Ham denoting the Hamilton model.

Figure 17 illustrates that there is a close association between the timing of switches to the high growth regimes by total GDP growth and by growth in the Services and Manufacturing sectors and to a lesser extent the Construction sector. The Primary and the Government and Community Services sectors display very different regime switching tendencies compared to total GDP and the other three sectors. The lack of any obvious relationship between the timing of Government and Community Services sector regime switching and GDP regime switching is consistent with earlier research by Kim, Buckle and Hall (1994) who found, using growth cycle methods, no significant contemporaneous correlation between cycles in real GDP and cycles in real government spending. The lack of any obvious relationship between the timing of Primary sector regime switching and GDP regime switching is more surprising.

The Manufacturing sector switches to the high state before the other sectors and total GDP during the 1981-1982 and the 1992–1996 high growth regimes. It also switches back to the low growth regime ahead of Services and ahead of GDP after the 1992–1996 high growth regime. However, Manufacturing does not always switch first. The Services sector was the first sector to transit to the 1999:1–2000:2 high growth regime and was the first to switch back to the low regime. Construction has more high growth regimes than the other sectors and was the last sector to switch to the low growth regime after mid 1990s high growth regime. It switched about one year after GDP and two years after Manufacturing switched back to a low growth regime.

There is considerable evidence that growth in many developed economies has become less volatile since the 1980s. McConnell and Perez–Quiros (1999) and Kim and Nelson (1999) provide compelling evidence of a sustained switch to lower volatility in US real GDP growth in early 1984. Their conclusions are based on applications of hidden Markov switching models. Using deviations from a moving average measure of trend growth, Blanchard and Simon (2001)

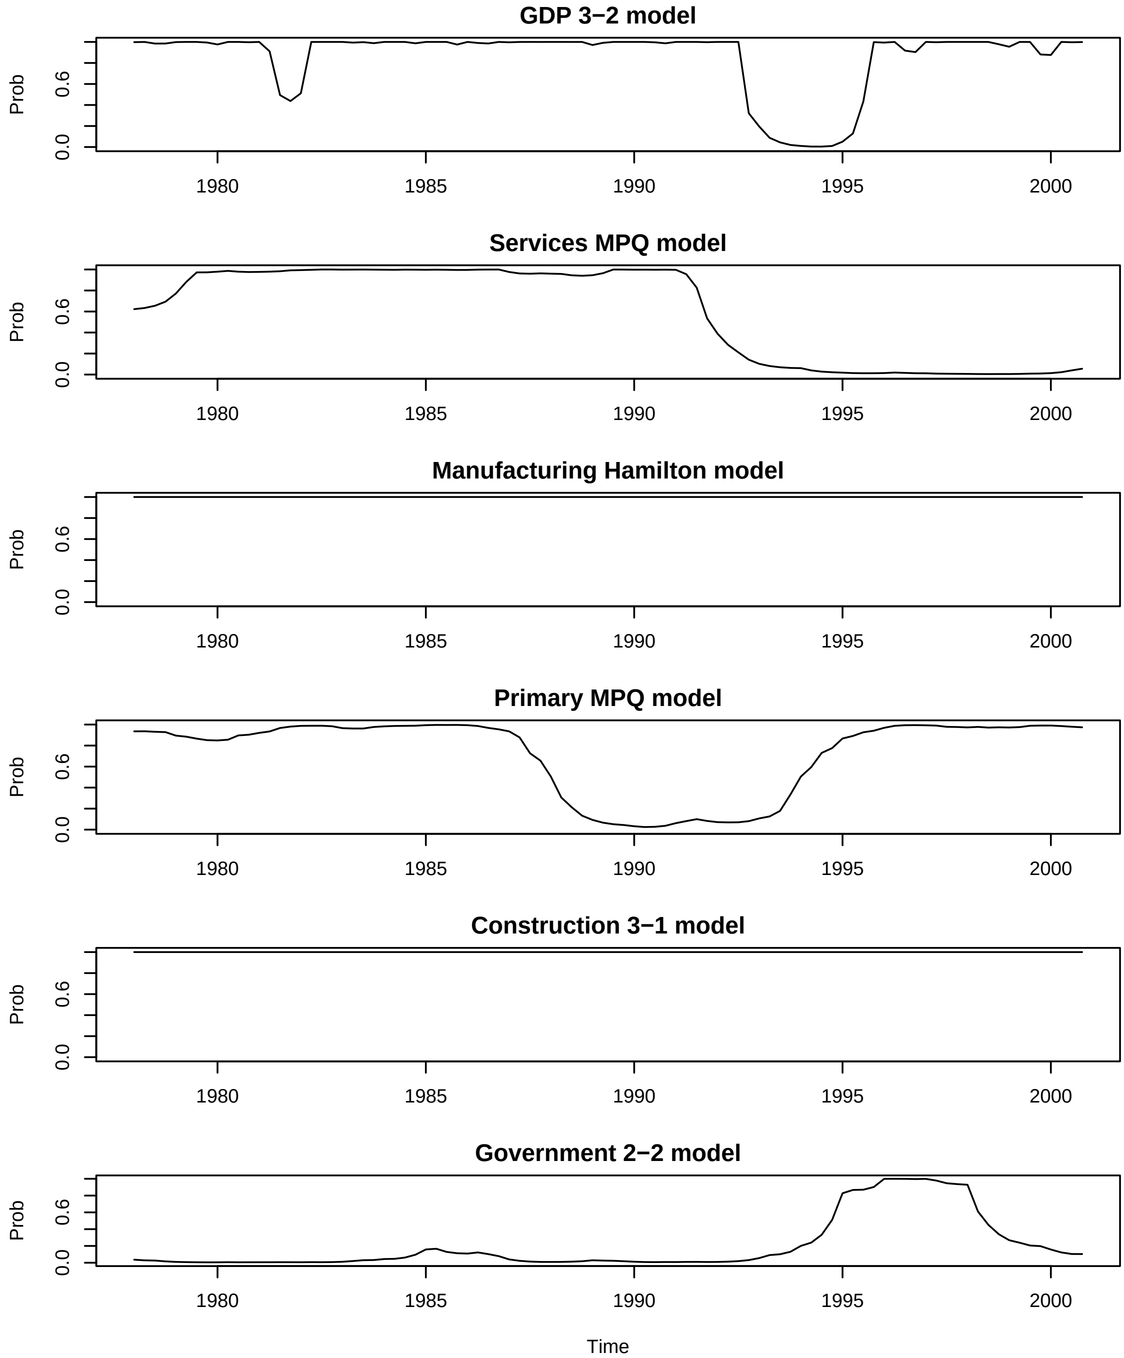

Figure 18: Probability of each production sector and total GDP growth rates being in a high volatility regime for each quarter from 1978:1 to 2000:4.

provide evidence of similar declines in Canada, United Kingdom, France, Germany and Simon (2001) for Australia.

Figure 18 shows the probability of New Zealand GDP and the five production sectors being in the high volatility regime over the period 1978:1 to 2000:4. New Zealand GDP growth is continuously in the high volatility regime throughout the late 1970s and the 1980s, but switches to the low volatility regime during the early to mid 1990s and remains in that state until 1996. However, in contrast to the overseas findings cited in the previous section, this switch to the low volatility regime is temporary and GDP has remained in the high volatility regime throughout the period since 1996.

The only production sector that switches to lower volatility in the early 1990s is the Services sector. This switch occurred in 1992:1, slightly earlier than the switch to lower volatility by total GDP. Moreover, in contrast to GDP and all other sectors, the decline in Services sector volatility is a sustained change to lower volatility. Nevertheless, despite contributing approximately 50 percent of total GDP, this is not sufficient to permanently change GDP volatility. This is because GDP volatility is affected by the covariance between the sectors growth rates which has a larger influence on overall GDP volatility than the volatility of the Services sector. Buckle, Haugh and Thomson (2001) show that a low covariance between sector growth rates was the main cause of the temporary decline in the GDP volatility during the 1990s. Furthermore, the switch to low volatility in the Services sector is also partially offset by switches to high volatility in the Primary sector and Government and Community Services sector. There are no volatility switches in the Manufacturing and Construction sectors as constant standard deviation models have been fitted to these sectors.

7 Summary and conclusions

This paper has reviewed and applied methodology for fitting hidden Markov switching models to New Zealand GDP data. Some key properties of the models have been developed together with a description of the estimation methods used, including maximum likelihood and the EM algorithm. Following a preliminary analysis of the GDP and production sector growth rates our findings are as follows.

Hidden Markov regime switching models provide a very flexible class of models that encompass a wide range of stationary time series behaviour from persistent cycles through to outliers. In this sense they are useful in their own right or as an exploratory tool to identify time series structure prior to fitting more appropriate dynamic models. They are relatively easy to fit to data using a mix of maximum likelihood and the EM algorithm.

The model has a conceptually simple structure with dynamics that operate on two time scales; the original quarterly time scale as well as the longer time scale embodied in the regimes. This simple, open framework allows the economic analyst to more easily interact with the data to allow enhanced economic and policy analysis.

On the whole, the HMM methods produce sensible results. When HMM trends and volatilities were compared to simple moving average estimates, there was good agreement. However the HMM trends and volatilities, in addition to being available at all time points, were much more adaptive and could track more abrupt changes in mean growth rates and volatilities by comparison to the simple moving average methods. Their potential to quickly adapt by comparison to other procedures means that they have the potential to provide useful and timely forecasts of future trend and volatility movements.

As found in many other studies, the hidden Markov switching model shows real potential for dating growth and volatility cycles. Indeed this is among the strongest of its attributes. However a potential weakness in this regard is the model's propensity to allow state visits of exactly one quarter. This is a consequence of the Markov chain assumption which means that the time spent in any state follows a geometric distribution which has a mode of unity. In particular this implies that any simulated HMM will show realisations where a state is visited for only one quarter. Figure 2 illustrates this in the case of a simulated MPQ model using parameter estimates fitted to US GNP growth rates. Solutions to this problem would be to modify the Markov chain appropriately or incorporate hidden semi–Markov models (see Ferguson (1980) or Sansom and Thomson (2001)).

Despite ongoing concerns about New Zealand's modest growth rate, the HMM models fitted to New Zealand's real GDP indicate New Zealand was in a high growth regime for an unusually long period of time during the early and mid 1990s. In general the hidden Markov models fitted to the production sectors revealed a close coherence between regime switching by GDP growth with the Services and Manufacturing sectors and to a lesser extent the Construction sector. However the mid 1990s was unusual in that all sectors simultaneously switched to high growth regimes during the early and mid 1990s. Declines in Manufacturing and Services sector growth rates appear to have led the switch of GDP to a low growth regime during 1996, followed much later by Construction. The Primary sector and Government and Community Services sectors remained in high growth regimes for the remainder of the decade.

Moreover, the mid 1990s experienced an unusually long period of low volatility. Switches to low volatility in the increasingly significant Services sector and the normally highly volatile Primary sector contributed to the decline in GDP volatility in the early and mid 1990s. A return to high volatility in the Primary sector and rising volatility in the Government and Community Services sector contributed to GDP switching back to the high volatility regime after 1995, despite sustained low volatility in the Services sector.

However these are preliminary findings only and more work is needed to fully explore the models and their variants, including suitable outlier models which may prove useful for the shorter noisy series experienced by small open economies such as New Zealand. Simulation studies would also help to build confidence in this technology by calibrating the accuracy and reliability of the estimation method and its out of sample performance in terms of prediction.

Finally and as alluded to in Section 6, the vector Markov switching models of Kontolemis (2001) and Krolzig (1997) would be an important direction for future development of these techniques in a New Zealand context. Such techniques hold out the promise of being more sensitive and timely tools for categorising sector cycles, identifying common cycles and exploring the correlation structure of total GDP and its production sectors.

Appendix: Recursions and likelihood

The probabilities $\gamma_t(j) = P(S_t = j|\mathbf{Y}, \boldsymbol{\theta}_0)$ and $\gamma_t(j,k) = P(S_t = j, S_{t+1} = k|\mathbf{Y}, \boldsymbol{\theta}_0)$ ($j, k = 1, \ldots, N$) need to be determined in order to evaluate the EM criterion $Q(\boldsymbol{\theta}|\boldsymbol{\theta}_0)$. However, as noted earlier, these probabilities are key quantities in their own right since they can be used to determine the likely states of the busines and volatility cycles given the data. Furthermore, the exact value of the likelihood is an important by–product of the recursive calculation procedures for the $\gamma_t(j)$ and $\gamma_t(j,k)$ given below.

Let $\mathbf{Y}_t = (Y_1, \ldots, Y_t)$ denote the observations up to an including time t ($t = 1, \ldots, T$) and define the likelihood ratios

$$\kappa_t = f(Y_t|\mathbf{Y}_{t-1}) = f(\mathbf{Y}_t)/f(\mathbf{Y}_{t-1}) \qquad (t = 2, \ldots, T) \tag{18}$$

with $\kappa_1 = f(Y_1)$. Here $f(Y_t|\mathbf{Y}_{t-1})$ is the predictive density of Y_t given $\mathbf{Y}_{t-1}$, $f(\mathbf{Y}_t)$ is the density of $\mathbf{Y}_t$ and $f(Y_1)$ is the density of Y_1. Now the logarithm of the likelihood (12) is given by

$$\log L = \log f(\mathbf{Y}_T) = \sum_{t=1}^{T} \log \kappa_t \tag{19}$$

where $\kappa_1 = \sum_{j=1}^{N} \pi_j f(Y_1|S_1 = j)$ and

$$f(Y_1|S_1 = j) = \frac{\sqrt{1-\rho^2}}{\sigma_j\sqrt{2\pi}} \exp\{-\tfrac{1}{2}(1-\rho^2)\left(\frac{Y_1 - \mu_j}{\sigma_j}\right)^2\}. \tag{20}$$

We also define the quantities $\alpha_t(j)$ and $\beta_t(j)$ where

$$\alpha_t(j) = P(S_t = j|\mathbf{Y}_t) \qquad (t = 1, \ldots, T)$$

and

$$\beta_t(j) = f(Y_{t+1}, \ldots, Y_T|S_t = j, Y_t) f(\mathbf{Y}_t)/f(\mathbf{Y}_T) \qquad (t = 1, \ldots, T-1)$$

with $\beta_T(j) = 1$ for all j. The $\alpha_t(j)$, $\beta_t(j)$ are proportional to the so-called *forward* and *backward* probabilities used in the HMM literature. Since the recursions used to compute the forward and backward probabilities given below reduce to accumulations of products of often very small probabilities, rescaling these probabilities is necessary to ensure that the recursive computations do not result in underflow. We have avoided this issue here by adopting the above definitions for the $\alpha_t(j)$, $\beta_t(j)$ which are suitably scaled versions of the forward and backward probabilities.

We now turn to the computation of the $\gamma_t(j)$, $\gamma_t(j,k)$ and the likelihood $f(\mathbf{Y}_T)$. First note that

$$\begin{aligned}\gamma_t(j) &= P(S_t = j|\mathbf{Y}_T) \\ &= f(Y_{t+1}, \ldots, Y_T|S_t = j, \mathbf{Y}_t) P(S_t = j|\mathbf{Y}_t) f(\mathbf{Y}_t)/f(\mathbf{Y}_T) \\ &= \alpha_t(j)\beta_t(j)\end{aligned} \tag{21}$$

holds for all t. Furthermore, using similar conditioning arguments,

$$\begin{aligned} \alpha_t(k) &= \sum_{j=1}^{N} P(S_{t-1}=j, S_t=k|\mathbf{Y}_t) \\ &= \sum_{j=1}^{N} f(Y_t|S_t=k, S_{t-1}=j, Y_{t-1}) P_{jk} \alpha_{t-1}(j)/\kappa_t \qquad (t=2,\ldots,T) \end{aligned} \tag{22}$$

where $f(Y_t|S_t=k, S_{t-1}=j, \mathbf{Y}_{t-1})$ is a Gaussian density with mean

$$E(Y_t|S_t=k, S_{t-1}=j, \mathbf{Y}_{t-1}) = \mu_k + \rho \frac{\sigma_k}{\sigma_j}(Y_{t-1}-\mu_j)$$

and variance σ_k^2. Values for the $\alpha_t(j)$ are now determined from the forward recursions (22) with

$$\kappa_t = \sum_{j=1}^{N}\sum_{k=1}^{N} f(Y_t|S_t=k, S_{t-1}=j, Y_{t-1}) P_{jk} \alpha_{t-1}(j) \qquad (t=1,\ldots,T)$$

and

$$\alpha_1(j) = f(Y_1|S_1=j)\pi_j/\kappa_1 \qquad (j=1,\ldots,N)$$

where $f(Y_1|S_1=j)$ is given by (20). A similar development for the $\beta_t(j)$ yields the backward recursions

$$\begin{aligned} \beta_t(j) &= \sum_{k=1}^{N} f(S_{t+1}=k, Y_{t+1},\ldots,Y_T|S_t=j, Y_t) f(\mathbf{Y}_t)/f(\mathbf{Y}_T) \\ &= \sum_{k=1}^{N} \beta_{t+1}(k) f(Y_{t+1}|S_{t+1}=k, S_t=j, Y_t) P_{jk}/\kappa_{t+1} \quad (t=T-1,\ldots,1) \end{aligned} \tag{23}$$

with $\beta_T(j)=1$ for all j. Also

$$\begin{aligned} \gamma_t(j,k) &= P(S_t=j, S_{t+1}=k|\mathbf{Y}_T) \\ &= \beta_{t+1}(k) f(Y_{t+1}|S_{t+1}=k, S_t=j, Y_t) P_{jk} \alpha_t(j)/\kappa_{t+1} \end{aligned} \tag{24}$$

holds for $t=1,\ldots,T-1$. These computationally efficient recursions provide all the elements needed to determine the log–likelihood of the data from (19), the $\gamma_t(j)$ from (21) and the $\gamma_t(j,k)$ from (24).

Appendix: additional figures

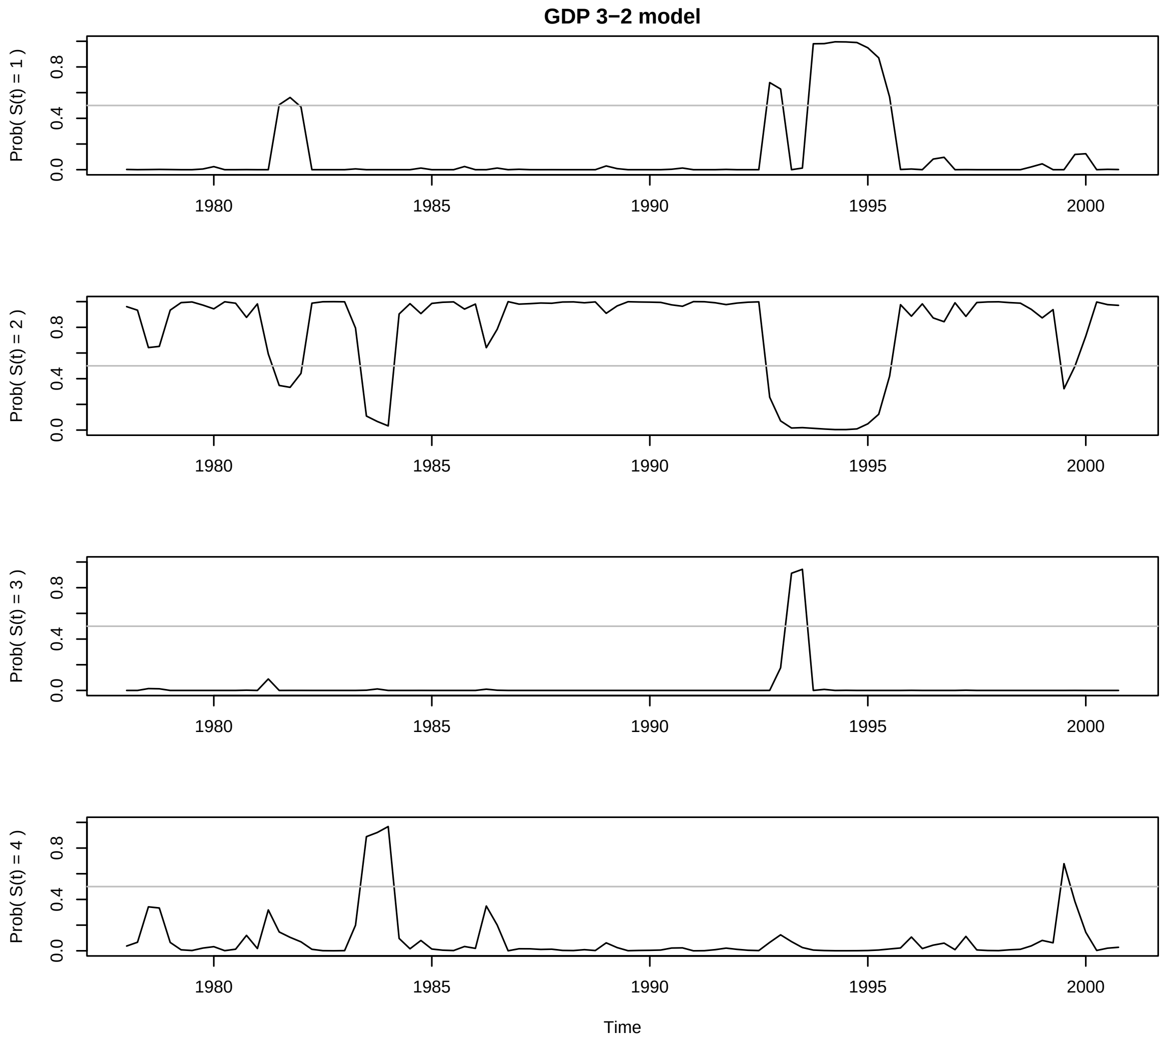

Figure 19: Plots of the $P(S_t = j|\mathbf{Y})$ for a 3–2 model fitted to quarterly GDP growth rates.

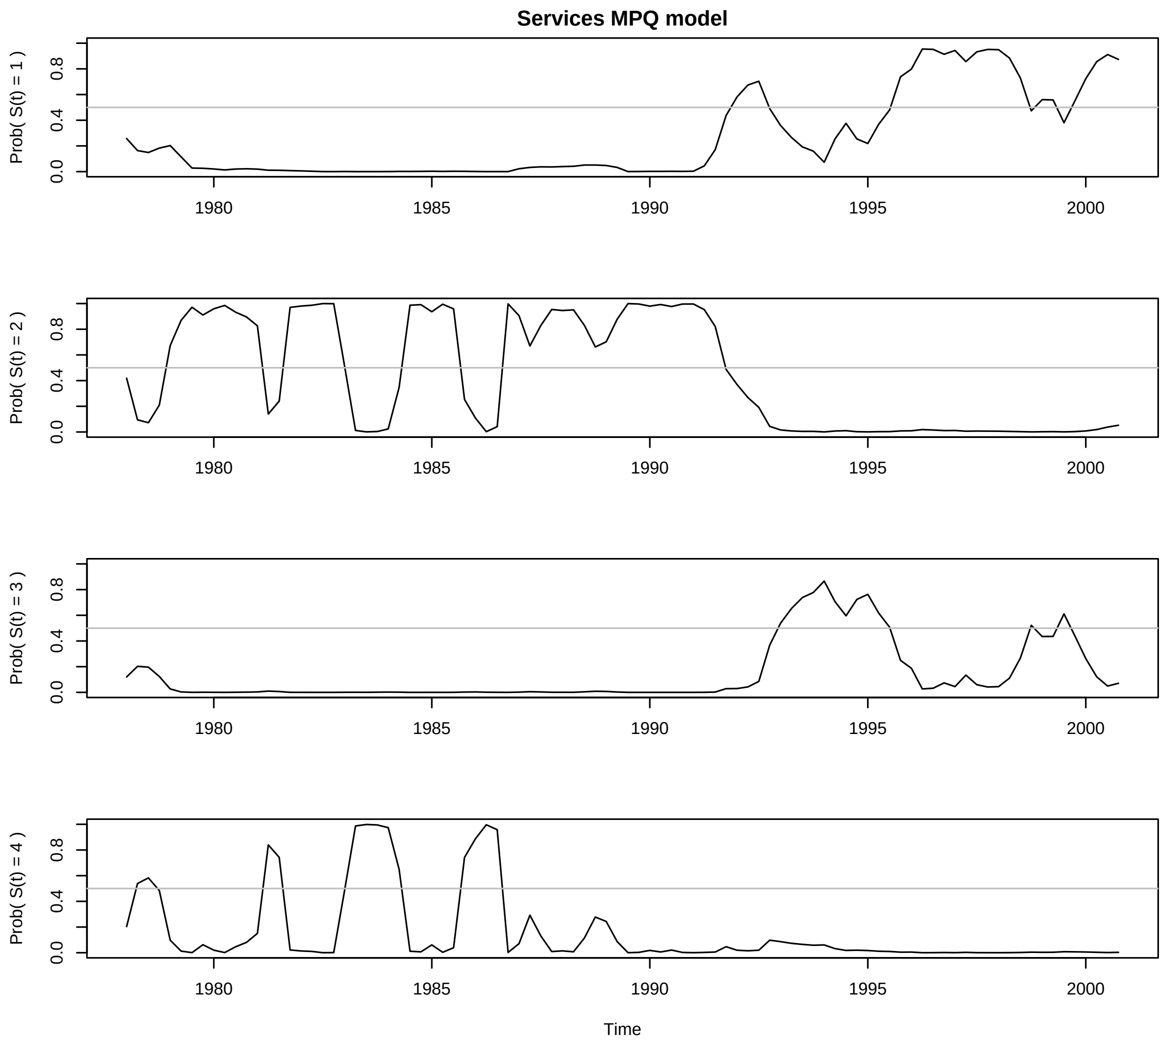

Figure 20: Plots of the $P(S_t = j|\mathbf{Y})$ for an MPQ model fitted to quarterly Services growth rates.

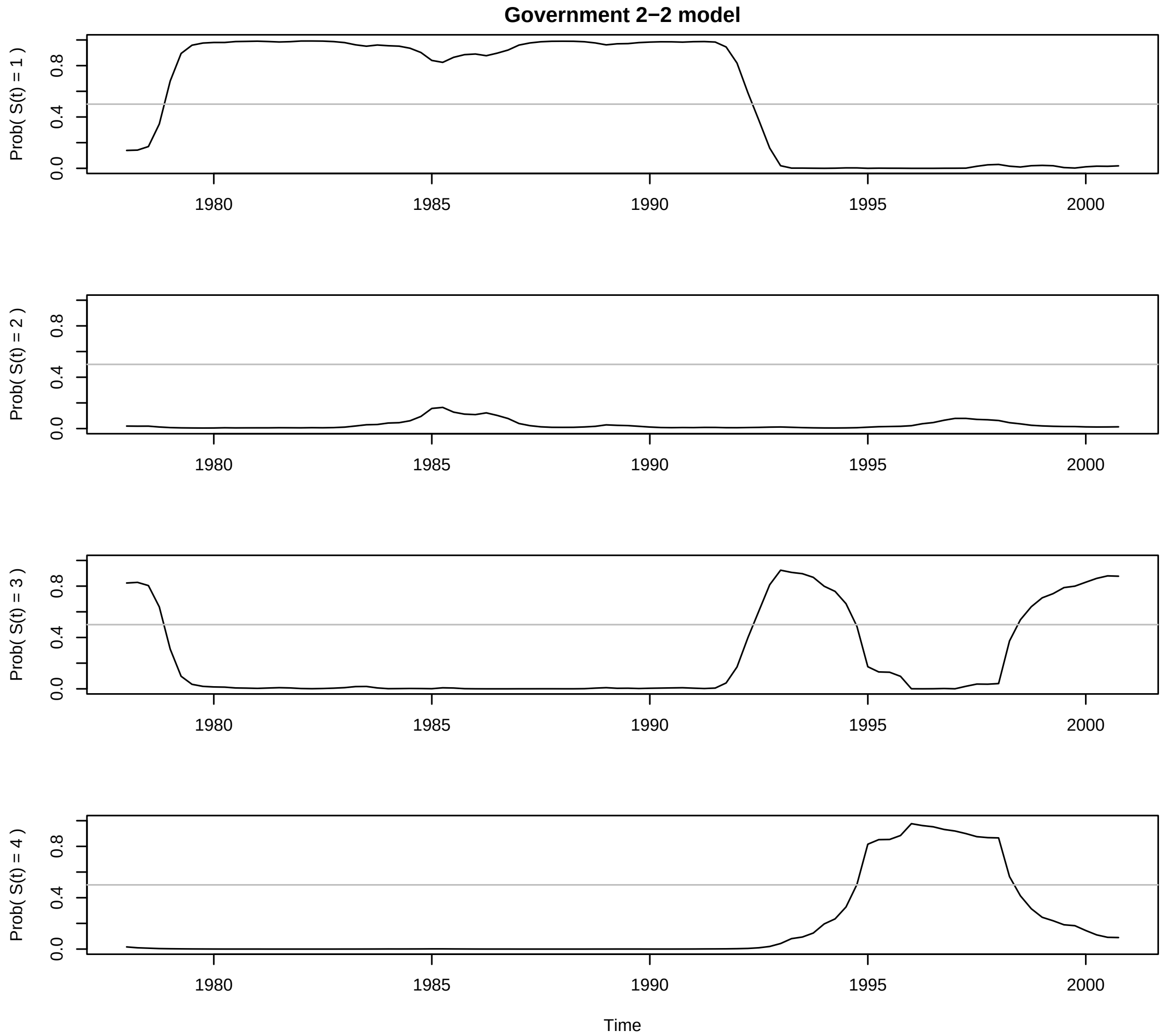

Figure 21: Plots of the $P(S_t = j|\mathbf{Y})$ for a 2–2 model fitted to quarterly Government and Community Services growth rates.

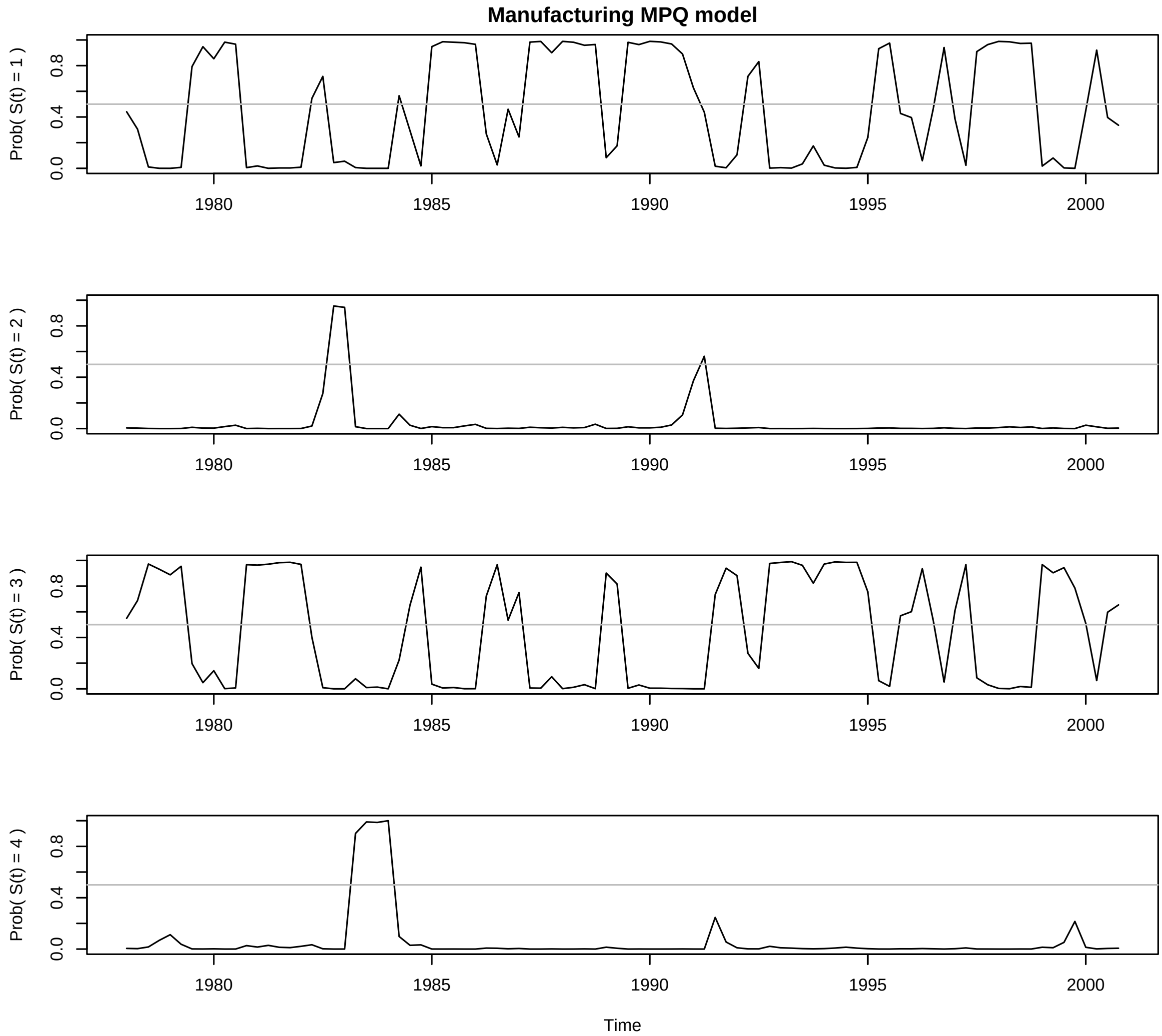

Figure 22: Plots of the $P(S_t = j|\mathbf{Y})$ for an MPQ model fitted to quarterly Manufacturing growth rates.

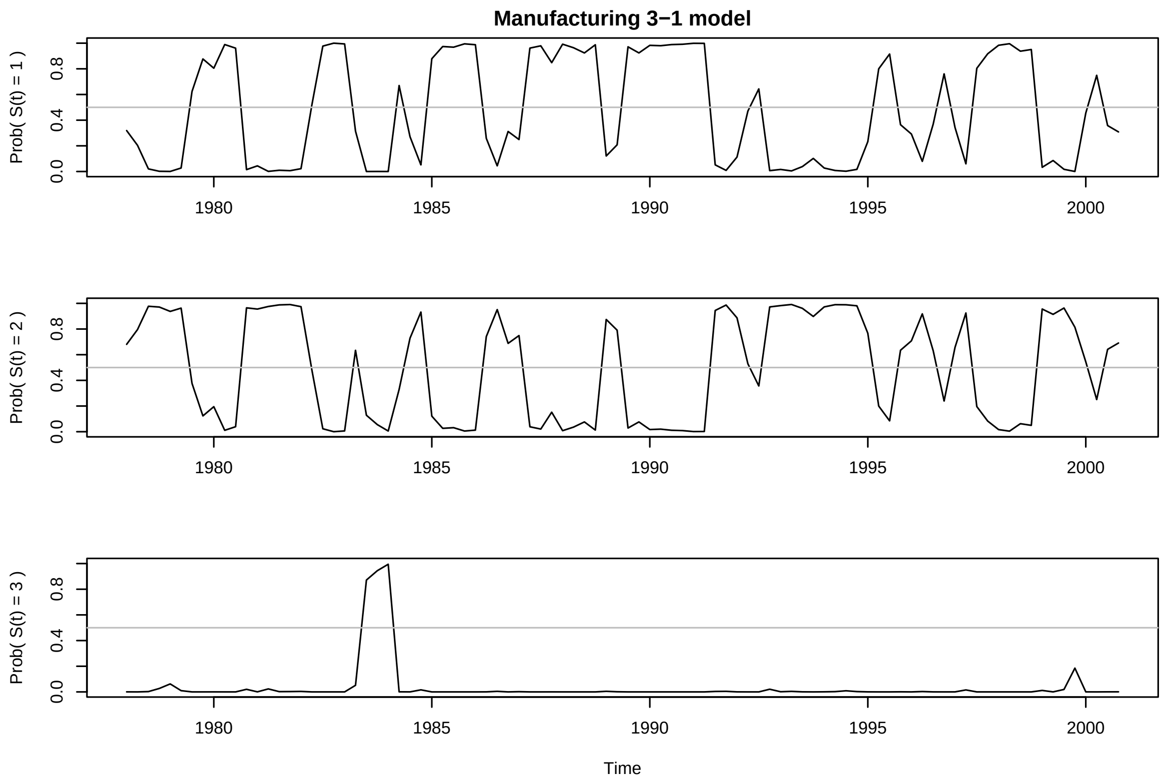

Figure 23: Plots of the $P(S_t = j|\mathbf{Y})$ for a 3–1 model fitted to quarterly Manufacturing growth rates.

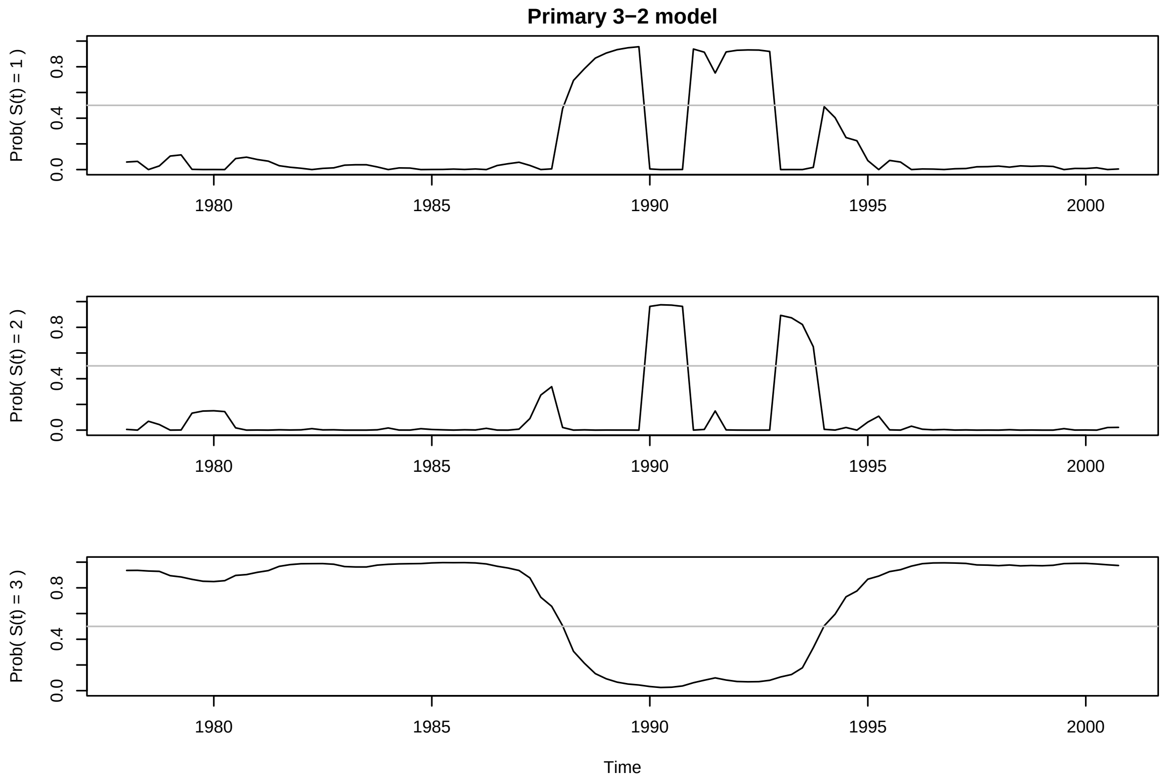

Figure 24: Plots of the $P(S_t = j|\mathbf{Y})$ for a 3–2 model fitted to quarterly Primary growth rates.

References

Baum, L.E. and Petrie, T. (1966) Statistical inference for probabilistic functions of finite state Markov chains. *Annals of Mathematical Statistics* **37**, 1554–1563.

Blanchard, O.J. and Simon, J. (2001) The long and large decline in U.S. output volatility. *Brookings Papers on Economic Activity* **1**.

Boldin, M.D. (1994) Dating turning points in the business cycle. *Journal of Business* **67**, 97–131.

Bry, G. and Boschan, C. (1971) Cyclical analysis of time series: Selected procedures and computer programmes. Technical paper No 20, National Bureau of Economic Research, New York.

Buckle, R.A., Haugh, D. and Thomson, P.J. (2001) Calm after the storm?: Supply-side contributions to New Zealand's GDP volatility decline. New Zealand Treasury Working paper 01/33. http://www.treasury.govt.nz/workingpapers/2001/01-33.asp

Buckle, R.A. and Carlson, J.A. (2000) Inflation and asymmetric price adjustment. *The Review of Economics and Statistics* **82**, 157–160.

Burns, A.F. and Mitchell, W.C. (1946) *Measuring business cycles*. National Bureau of Economic Research, New York.

Cover, J.P. (1992) Asymmetric effects of positive and negative money-supply shocks. *The Quarterly Journal of Economics* **107**, 1261–1282.

De Long, J.B and Summers, L.H. (1988) How does macroeconomic policy affect output? *Brookings Papers on Economic Activity* **2**, 433–480.

Dempster, A.P., Laird, N.M., and Rubin, D.B. (1977) Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B* **39**, 1–38.

Diebold, F.X. and Rudebusch, G.D. (1992) Have post–war economic fluctuations been stabilised? *American Economic Review* **83**, 993–1005.

Durland, M.J. and McCurdy, T.H. (1994) Duration dependent transitions in a Markov model of US GNP growth. *Journal of Business and Economic Statistics* **12**, 279–288.

Elliot, R.J., Aggoun, L., and Moore, J.B. (1995) *Hidden Markov models*. Springer-Verlag, New York.

Engle, C. and Hamilton, J.D. (1990) Long swings in the dollar: are they in the data and do markets know it? *American Economic Review* **80**, 689–713.

Ferguson, J.D. (1980) Variable duration models for speech. In *Symposium on the Application of hidden Markov models to text and speech* (ed. J.D. Ferguson), Institute for Defense Analyses, Princeton N.J., 143–179.

Filardo, A.J. (1994) Business cycle phases and their transitional dynamics. *Journal of Business and Economic Statistics* **12**, 299–308.

Goldfeld, S.M. and Quandt, R.E. (1973) A Markov model for switching regressions. *J. Econometrics* **1**, 3–16.

Hamilton, J.D. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrika* **57**, 357–384.

Hamilton, J.D. and Susmel,R (1994) Autoregressive conditional heteroskedasticity and changes in regime. *J. Econometrics* **64**, 307–333.

Harvey, A.C. (1989) *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Harvey, A.C. and Jaeger, A. (1993) Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* **8**, 231–247.

Haugh, D. (2001) Calibration of a chain volume production GDP database. New Zealand Treasury, MVD Technical Note 01/2001.

Hodrick, R.J. and Prescott, E.C. (1980) Postwar US business cycles: an empirical investigation. Discussion paper No. 451, Carnegie-Mellon University.

Karras, G. (1996) Are the output effects of monetary policy asymmetric? Evidence from a sample of European countries. *Oxford Bulletin of Economics and Statistics* **58**, 267–278.

Kim, CJ. (1994) Dynamic linear models with Markov–switching. *J. Econometrics* **60**, 1–22.

Kim, CJ and Nelson, C.R. (1999a) *State–space models with regime–switching: classical and Gibbs–sampling approaches with applications*. MIT Press, USA.

Kim, CJ and Nelson, C.R. (1999b) Has the U.S. economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle. *The Review of Economics and Statistics* **81**, 608–616.

Kim, KH, Buckle, R.A. and Hall, V. B, (1994) Key features of New Zealand business cycles, *The Economic Record* **70**, 56–72.

Kim, KH, Buckle, R.A. and Hall, V.B. (1995) Dating New Zealand business cycles. *New Zealand Economic Papers* **29**, 143–171.

Kontolemis, Z.G. (2001) Analysis of the US business cycle with a vector-Markov-switching model. *Journal of Forecasting* **20**, 47–61.

Koopman, S.J., Harvey, A.C., Doornik, J.A. and Shepherd, N. (1995) *STAMP 5.0: Structural Time Series Analyser, Modeller and Predicter*, Chapman and Hall, London.

Krolzig, HM (1997) *Markov–switching vector autoregressions: modelling, statistical inference, and application to business cycle analysis*. Lecture Notes in Economics and Mathematical Systems **454**, Springer–Verlag, New York.

Lam, PS (1990) The Hamilton model with a general autoregressive component. *Journal of Monetary Economics* **26**, 409–432.

Levinson, S.E., Rabiner, L.R., and Sondhi, M.M. (1983) An introduction to the application of the theory of probabilistic functions of a Markov process to automatic speech recognition. *Bell Systems Technical Journal* **62**, 1035–1074.

MacDonald, I.L. and Zucchini, W. (1997) *Hidden Markov and other models for discrete-valued time series*. Chapman and Hall, London.

McConnell, M.M and Perez-Quiros, G. (2000) Output fluctuations in the United States: What has changed since the early 1980's? *The American Economic Review* **90**, 1464–1476.

Meilijson, I. (1989) A fast improvement to the EM algorithm in its own terms. *Journal of the Royal Statistical Society B* **51**, 127–138.

Mitchell, W.C. (1913) *Business Cycles*. University of California Press.

Mitchell, W. C. (1927) *Business cycles: The problem and its setting*. National Bureau of Economic Research, New York.

Morgan, D.P. (1993) Asymmetric effects of monetary policy. *Federal Reserve Bank of Kansas City Economic Review* **78**, 21–33.

Neftci, S.N. (1984) Are economic time series asymmetric over the business cycle? *Journal of Political Economy* **92**, 307–328.

Olekalns, N. (1995) The effect of nominal demand shocks on manufacturing output: Evidence from disaggregated Australian data. *The Economic Record* **71**, 66–76.

Rabiner, L.R. (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* **77**, 257–285.

Rhee, W. and Rich, R.W. (1995) Inflation and the asymmetric effects of money on output fluctuations. *Journal of Macroeconomics* **17**, 683–702.

Sansom, J. and Thomson, P.J. (2001) Fitting hidden semi–Markov models to breakpoint rainfall data. Journal of Applied Probability **38A**, 142–157.

Shaghil, A., Levin, A. and Wilson, B.A. (2001) Recent U.S. macroeconomic stability: Good luck, good policies or good practices? Mimeo, The Board of Governors of the Federal Reserve System, Washington DC 20551.

Sichel, D.E. (1987) Business cycle asymmetry: A deeper look. Mimeo, Princeton University.

Simon, J. (2001) The decline in Australian output volatility. Reserve Bank of Australia, Research Discussion Paper RDP 2001–01.
http://www.rba.gov.au/PublicationsAndResearch/rdp/RDP2001–01

Zucchini, W. and Guttorp, P. (1991) A hidden Markov model for space-time precipitation. *Water Resources Research* **27**, 1917–1923.