

Testing for Regime Switching in Singaporean Business Cycles

Robert Breunig*[†]
School of Economics
Faculty of Economics and Commerce
Australian National University

and

Alison Stegman
Research School of Pacific and Asian Studies
Australian National University

September 27, 2003

Abstract

We examine a Markov Switching model of Singaporean GDP using a combination of formal moment-based tests and informal graphical tests. The tests confirm that the Markov Switching model fits the data better than a linear, autoregressive alternative. The methods are extended to allow us to identify precisely which features of the data are better captured by the non-linear model. The methods described here allow model selection to be related to the intended use of the model.

Keywords: *Markov Switching Models, Specification Testing, Nonparametric Estimation, Moment Tests*

JEL Classification: C12,C14,C22,C52

*Address for correspondence: School of Economics, Faculty of Economics and Commerce, Australian National University, Canberra ACT 0200, AUSTRALIA, email: Robert.Breunig@anu.edu.au

[†]We would like to thank Philip Bodman and Mark Crosby for their comments and Adrian Pagan for suggesting this data series as an interesting one for analysis. All errors are our own.

1 Introduction

Diagnostic testing forms an integral part of applied economic research. In linear modelling, testing procedures are well developed and easy to implement. The adaptation of such procedures to nonlinear modelling, however, is often complex and, as a result, researchers undertaking complicated nonlinear modelling often provide little diagnostic evidence to support their results. This represents a significant impediment to the development of new ideas and the discussion of controversial results. Without diagnostic results, readers are unable to form conclusions regarding the validity and significance of researchers' conclusions, and researchers run the risk of having potentially important results overlooked.

The purpose of this paper is to introduce readers of this journal to some recent developments in the testing of Markov-Switching models. These tests are applied to the Bodman and Crosby (1999) Markov-Switching model of Singaporean Gross Domestic Product (GDP) which appeared previously in this review. We agree with their conclusion that the Markov-Switching model outperforms a linear autoregressive model. We present several tests that support the Markov-Switching model and we show graphically exactly which aspects of the data are better represented by the non-linear model. It is hoped that this study will persuade other researchers to include a similar analysis when reporting their results.

The following section briefly reviews the theory of Markov-Switching models and introduces the testing procedures. Section three contains the main results of the tests and our conclusions about model selection. We conclude in the final section.

2 Model and Testing Procedures

2.1 Markov-Switching Models

Many economic variables behave differently depending upon the stage of the business cycle. Researchers have attempted to account for such asymmetries by modelling the transition between a ‘high growth’ phase and a ‘low growth’ phase as a regime switch. Markov-Switching (MS) models are particularly appealing because they allow this switching process to be endogenized and allow for inferences regarding the timing and nature of such switches. The seminal paper of Hamilton (1989), in which he models United States GDP as a two-state Markov-Switching process, encouraged the use of these techniques and led to the wide application of MS models in applied economic research. As in Hamilton (1989), this study examines GDP behavior. GDP growth is allowed to behave differently, depending on whether the economy is in a high growth phase or a low growth phase.

The model is characterized as follows: A state variable at time t , s_t , describes the state of the economy. When $s_t = 0$, the economy is in a low growth phase and when $s_t = 1$ the economy is in a high growth phase. GDP growth is modelled as

$$y_t = \alpha_0 + \alpha_1 s_t + \sum_{p=1}^P \phi_p (y_{t-p} - \alpha_0 - \alpha_1 s_{t-p}) + \sigma_s \varepsilon_t \quad (1)$$

where α_0 is the mean in state 0 (low growth), $\alpha_0 + \alpha_1$ is the mean in state 1 (high growth), ϕ_p are the autoregressive coefficients, σ_0 is the standard deviation in state 0, σ_1 is the standard deviation in state 1, and ε_t is a white noise error process.

The state variable s_t is unobservable and assumed to follow a first order Markov chain:

$$p_{ij} = P(s_t = j | s_{t-1} = i, s_{t-2} = k, \dots) = P(s_t = j | s_{t-1} = i) \quad (2)$$

The probability of being in a given state depends on the past only through the most recent value s_{t-1} . p_{ij} gives the transition probability that state i will be followed by state j . These models are generally estimated by maximum likelihood, although Bayesian estimation of MS models is gaining popularity.

2.2 Testing procedures

Problems associated with MS models (and nonlinear models in general) can be characterized into two main areas: convergence and specification. We follow the testing procedures developed in Breunig et al. (2003), which are aimed at both convergence and specification problems and are designed to be simple enough to enable use in the “discovery process”. Wald and Likelihood Ratio tests (variants of which have been proposed for MS models by Garcia (1998) and Hamilton (1996)) can be difficult to implement in nonlinear models and this perhaps explains why their use is so infrequent. The tests used here have the advantage of being simple to implement and interpret.

This section presents both the formal and informal procedures proposed by Breunig et al. (2003). The procedures are based on a comparison of the ‘sample’ properties of the data with the ‘population’ characteristics suggested by the model and have general applicability to any non-linear problem. First, the hypothesized model is estimated by maximum likelihood. These coefficients (and the assumed process) are then used to simulate a large set of pseudo-

observations. It is assumed that this set is large enough so that there is no error (or variance) attached to the simulation process. These observations are then interpreted as the ‘population’ implied by the estimated model.

The formal tests involve statistically comparing the “sample” characteristics of the data with the “population” characteristics of the data in a moment based test. The null hypothesis is that the given data set is a sample from the pseudo population (here implying an underlying MS model consistent with the estimated parameters and the distributional assumptions of the estimation technique), whilst the alternative is that it is not.

Formally, if $\hat{\mu}$ is the estimated mean of some function of the data and $\hat{\theta}$ are maximum likelihood estimates associated with the MS model, then $\mu(\hat{\theta})$ is the associated mean implied by the model and the null and alternative hypotheses are given by

$$H_0 : \mu = \mu(\theta_0)$$

$$H_1 : \mu \neq \mu(\theta_0)$$

where θ_0 is the true value of θ under the null hypothesis that the MS model specification is correct. Acceptance of the null hypothesis implies that the process that generated the “population” was ‘valid’. This can be thought of as a parametric encompassing test.

The test statistic is a Hausman (1978)-type test which compares one estimator which is consistent ($\hat{\mu}$) to another which is only consistent under the null, but is more efficient when the null is correct. It is formed as

$$\hat{\tau} = \left(\hat{\mu} - \mu(\hat{\theta}) \right)' \left[\text{var} \left(\hat{\mu} - \mu(\hat{\theta}) \right) \right]^{-1} \left(\hat{\mu} - \mu(\hat{\theta}) \right) \quad (3)$$

The variance of this statistic is difficult to calculate and Breunig et al. (2003) suggest replacing $var(\hat{\mu} - \mu(\hat{\theta}))$ with $var(\hat{\mu})$, which provides a conservative test. See their paper for the technical details, as well as for information on how to compute a variance which would take into account the variability arising from the estimate of the MS model. In the results presented below, we use a robust estimator of $var(\hat{\mu})$, calculated using the Newey-West formula with the lag set equal to 9.

The test statistic is constructed for four functions of the data: the mean, the variance, the probability of observing a contraction after an expansionary period, and the probability of observing an expansion after an expansionary period. These two latter quantities are estimated by

$$\hat{p}_1 = \frac{1}{T} \sum_{t=1}^T 1(y_{t-1} < 0, y_t > 0) \quad (4)$$

and

$$\hat{p}_2 = \frac{1}{T} \sum_{t=1}^T 1(y_{t-1} > 0, y_t > 0) \quad (5)$$

where $1(\cdot)$ is the indicator function.

The motivation for the mean and variance tests is that a good model should be able to capture the basic moments of the data. The motivation for the ‘quadrant’ tests is that a model of GDP growth should provide accurate predictions of the probability of expansions and contractions for forecasting purposes.

As an informal test, Breunig et al. (2003) suggest plotting non-parametric densities and conditional means of the ‘pseudo population’ against a scatter plot of the actual data to provide an indication of how well the fitted model is able to capture non-linearities in the data. In this case, we will plot the conditional

mean $E(y_t|y_{t-1})$ which is of interest for forecasting. This also will provide information about the quadrant tests (4) and (5). The conditional mean is plotted using standard Nadaraya-Watson kernel regression.¹

3 Data, Estimates, and Test Results

3.1 Data

To reproduce the MS model of Bodman and Crosby (1999) (hereafter referred to as BC) we use quarterly observations of Singaporean GDP over the period 1975:1 to 1999:2. The data was obtained from the Singapore Department of Statistics. Non-seasonally adjusted data were available for the entire period, seasonally adjusted observations were only available for the period 1978:1 to 1999:2. The 12 observations from 1975 to 1977 were therefore adjusted using the seasonal factors for 1978. The results presented below are not affected by this assumption². Singaporean GDP is non-stationary, so we use log differences (GDP growth) as the data series for analysis.³

¹See Pagan and Ullah (1999) for details. We use a range of bandwidths and kernel weighting functions. As we use 60,000 simulated observations, it is not particularly surprising that the results are not sensitive to these choices.

²We used seasonally adjusted data and the data range from 1975:1 to 1999:2 in an attempt to match Bodman and Crosby (1999) as closely as possible. Bodman and Crosby provided us with their original data set, but we chose instead to use the data currently available from the Singapore Department of Statistics. It is worth noting that the Singapore Department of Statistics has recently updated and re-issued the GDP series. If we attempt to date the business cycle using the revised data we get slightly different turning points than using the original data of Bodman and Crosby (1999). This is not important for our paper, however, as the MS models implied by the two data sets are almost identical as shown below. The results presented in this study are robust to a range of seasonal adjustment assumptions and also robust to dropping the pre-1978 observations.

³Details of unit root tests and graphs of raw and transformed data, as well as the data itself, are available from the authors.

3.2 Estimates

Standard Box-Jenkins analysis reveals that an autoregressive model with 2 lags (AR(2)) is the appropriate specification for a linear model. This is also the specification chosen by using the minimum Akaike Information Criterion. Table 1 provides the AR(4) estimates of BC along with our AR(4) and AR(2) estimates for comparison. Following BC, we apply the BDS test to the residuals from the various models. We use the small sample significance levels of Kanzler (1999). See Brock et al. (1996) for details of the test.

Table 1: Linear Model Estimates

$$y_t = \delta_0 + \sum_{l=1}^L \delta_l y_{t-l} + \sigma \varepsilon_t$$

	Bodman and Crosby (1999), AR(4)	AR(4)	AR(2)
δ_0	1.25* (0.33)	1.87* (0.22)	1.87* (0.23)
δ_1	0.20 (0.11)	0.15 (0.10)	0.15 (0.10)
δ_2	0.20 (0.11)	0.21* (0.10)	0.20 (0.10)
δ_3	0.15 (0.11)	0.08 (0.10)	
δ_4	-0.16 (0.11)	-0.13 (0.12)	
σ	1.40* (0.24)	1.36* (0.09)	1.37* (0.09)
BDS (m=2)	2.46*	2.75*	3.58*

*indicates significant at 5% level

For both the AR(2) and the AR(4) models, the BDS test rejects the null hypothesis of independently and identically distributed residuals. BC draw a similar conclusion for their model and argue that this provides evidence of nonlinearities in the data. They estimate a two state MS model with four

autoregressive lags (MS(4)) in an attempt to account for these nonlinearities. The results of their estimation are presented in the second column of Table 2. We replicate these results by using the Gauss estimation code available on Hamilton's web site (<http://econ.ucsd.edu/~jhamilto/>) and the revised data from the Singapore government. The third column of Table 2 reports the results of our re-estimation of the MS(4) model. We also estimate the MS model with 2 lags (MS(2)) (the appropriate specification suggested by our linear ARIMA analysis). These results are shown in the last column of Table 2.

Table 2: MS Model Estimates

	MS(4)		MS(2)
	Bodman and Crosby (1999)	Our estimates	
μ_0	-0.854* (0.085)	-0.960* (0.269)	-1.173* (0.461)
μ_1	3.090* (0.488)	3.158* (0.351)	3.287* (0.485)
ϕ_1	-0.274* (0.108)	-0.283* (0.119)	-0.256* (0.116)
ϕ_2	0.059 (0.106)	0.047 (0.114)	-0.050 (0.121)
ϕ_3	0.309* (0.094)	0.271* (0.109)	
ϕ_4	0.283* (0.109)	0.227 (0.123)	
σ_0	0.482* (0.258)	0.841 (0.654)	1.076 (0.875)
σ_1	1.109* (0.199)	1.046* (0.419)	1.116* (0.443)
p_{00}	0.67	0.67* (0.15)	0.69* (0.17)
p_{11}	0.96	0.96* (0.02)	0.97* (0.02)
log-likelihood	-152	-151.6	-157.4
BDS (m=2)		0.0604	-0.4745

A comparison of the re-estimated results with those presented in BC suggests no serious convergence problems. They also suggest that while the data revision undertaken by the Singapore government affects the dating of the business cycle, model estimates based upon the revised data are not affected. All of the re-estimated coefficients lie within a 95 percent confidence interval of those reported by BC. The re-estimated results suggest that the mean growth rate of Singaporean GDP in a low growth state is -0.96 , whilst the mean in a high growth state is $3.16 - 0.96 = 2.2$. The business cycle also shows evidence of asymmetry with the probability of leaving a high growth state, 0.04 , significantly lower than the probability of leaving a low growth state, 0.33 .

The results from the MS(2) model are similar to the results from the MS(4) model. BC conduct the BDS test on the residuals from their MS(4) model. Presumably, they conduct their diagnostic analysis on this alternative model to simplify the calculation and interpretation of the test statistics (they also conduct a Likelihood Ratio test to compare this constant variance model with the linear AR(4) model and reject the null hypothesis of a constant mean across states). We therefore use the BDS test on the residuals from our re-estimated MS(4) and MS(2) models. We calculate the residuals from the MS models by constructing a weighted sum of the 32 possible residuals for each time period. The weights for each time period are the conditional probabilities corresponding to each state. The results of this test are presented in the last row of Table 2. In both cases, we were unable to reject the null hypothesis of independently and identically distributed residuals. We are unsure, however, how to interpret this result given that each of the residuals has been constructed as a weighted sum of all of the possible residuals given the economy is in either state 0 or

state 1 in any given time period. We are also unsure of the behavior of the test statistic under these conditions. We argue that the testing procedures of section 2.2 provide a more intuitive set of diagnostic checks.

3.3 Test Results

Table 3 summarizes the moments that we find in the data as well as the moments implied by the five estimated models: MSBC refers to the estimates of BC and the other models are estimated using the revised Singaporean GDP data.

Table 3: Data and Model Moments

	Data	MSBC	MS(4)	MS(2)	AR(4)	AR(2)
μ	1.84	1.90	1.91	1.79	2.69	2.75
σ^2	2.01	2.29	1.95	2.22	1.08	1.07
\hat{p}_1	0.05	0.06	0.06	0.06	0.01	0.00
\hat{p}_2	0.85	0.81	0.85	0.83	0.99	0.99

We conduct the moment tests of section 2 on these five models and report the test results in Table 4. Rejection of the moment test implies a rejection of the model.

Table 4: Moment tests

	MSBC	MS(4)	MS(2)	AR(4)	AR(2)
$\hat{\tau}(\mu)$	-0.51	-0.38	0.26	-4.54*	-4.82*
$\hat{\tau}(\sigma^2)$	-0.67	0.13	-0.49	2.19*	2.224*
$\hat{\tau}(\hat{p}_1)$	-0.63	-0.21	-0.61	2.06*	2.37*
$\hat{\tau}(\hat{p}_2)$	0.45	-0.09	0.36	-2.43*	-2.43*

The Markov Switching models capture the important features of the data

very well. All of the test statistics for the nonlinear models are insignificant, and although this is a conservative test, all of the statistics are well below the significance level of 1.96. The linear models, in contrast, perform badly in all categories.

Figure 1 plots the non-parametric mean of y_t conditional on y_{t-1} . MSBC denotes the conditional mean implied by the model estimated by BC, MS4 denotes the conditional mean implied by the re-estimation of the Markov Switching model with four autoregressive lags, and L4 is the conditional mean implied by the corresponding linear AR(4) model. Theoretically this last line should be perfectly linear but outliers at the extremes of the plot area act to distort the estimation slightly. Such boundary problems also occur in estimating the conditional mean for the MS models and conclusions should be based primarily on the area where the observations are concentrated. The graph not only suggests that nonlinearities are present, but it provides detailed information on the type of non-linearity implied by the MS model. The MS model captures three main features of the data that are missed by the linear model:

1. The low expected value for the growth rate in periods following negative growth.
2. The sharp increase in the expected growth for next period when current period growth rate is between -1% and 1%
3. The decline in the expected value of next period's growth rate when current growth is over 1%.

The plot of the MS(2) model is similar and thus is not shown here.

As a general conclusion, both of the MS models appear to represent an advance over their corresponding linear models and we would expect researchers

interested in forecasting to obtain superior results by utilizing the information contained in these models. We note that non-linear models have been criticized in general for failing to improve out-of-sample forecasts relative to linear models. van Dijk and Franses (2003) suggest alternative forecast evaluation criteria using weights that could be in part decided by the state probabilities.

Figure 1 here

4 Conclusion

We have presented a set of moment tests which applied researchers can use to test whether an estimated non-linear model is able to match moments of the data. Moments can be chosen according to the particular question being addressed. Here, we use the mean, the variance, and information about one-step ahead predictions to test a model of Singaporean GDP growth. This model was previously proposed by Bodman and Crosby (1999) and this paper finds evidence to support their conclusion that their Markov-Switching model is superior to the linear alternative. We extend their conclusion by showing exactly which features of the data are best described by the Markov Switching model.

The results presented here contrast with those of Breunig et al. (2003) and Breunig and Pagan (2003), where moment tests and graphical methods were used to cast doubt on the validity of Markov-Switching models. The example presented here provides a nice illustration of how the tests can be used by researchers to reinforce conclusions in favor of Markov-switching models.

We also use graphical methods to explore the added value of the non-linear model. In this case, the Markov-switching model nicely captures several important features of the data. None of this information is available in the standard

analysis which is usually presented with estimates of Markov-switching models. Our hope is that applied researchers will use this combination of formal moment tests and informal graphical tests in future research.

References

- Bodman, P. and Crosby, M. (1999). Non-linearities in the Singaporean business cycle. *Singapore Economic Review*, 44(2):24–35.
- Breunig, R. and Pagan, A. (2003). Do Markov-switching models capture nonlinearities in the data? Tests using nonparametric methods. *Mathematics and Computers in Simulation*. Forthcoming.
- Breunig, R. V., Najarian, S., and Pagan, A. R. (2003). Specification testing of Markov-switching models. *Oxford Bulletin of Economics and Statistics*. Forthcoming.
- Brock, W., Dechert, D., Scheinkman, J., and LeBaron, B. (1996). A test for independence based on the correlation dimension. *Econometric Reviews*, 15(13).
- Garcia, R. (1998). Asymptotic null distribution of the likelihood ratio test in markov switching models. *International Economic Review*, 39:763–788.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384.
- Hamilton, J. D. (1996). Specification testing in Markov-switching time-series models. *Journal of Econometrics*, 70:127–157.
- Hausman, J. (1978). Specification tests in econometrics. *Econometrica*, 46:1251–1271.
- Kanzler, L. (1999). Very fast and correctly sized estimation of the BDS test. Available at <http://www2.gol.com/users/kanzler/>.
- Pagan, A. R. and Ullah, A. (1999). *Nonparametric Econometrics*. Cambridge University Press.
- van Dijk, D. and Franses, P. H. (2003). Selecting a nonlinear time series model using weighted tests of equal forecast accuracy. Econometric Institute Report EI 2003-10.