Regime Switches in GDP Growth and Volatility: Some International Evidence and Implications for Modelling Business Cycles*

Penelope A. Smith and Peter M. Summers
Melbourne Institute of Applied Economic and Social Research
The University of Melbourne

Melbourne Institute Working Paper No. 21/02

ISSN 1328-4991 (Print)

ISSN 1447-5863 (Online)

ISBN 0 7340 3105 X

November 2002

*A previous version of this paper circulated with the title "On the interactions between growth and volatility in a Markov switching model of GDP." We are grateful to Don Harding, Adrian Pagan, and participants at the 22nd International Symposium on Forecasting in Dublin and the 2002 Econometric Society Australasian Meeting in Brisbane for helpful comments and discussions. Computations in this paper were carried out (in part) using the Bayesian Analysis, Computation and Communications (BACC) software available at www.econ.umn.edu/~bacc. We retain all responsibility for any remaining errors.

Melbourne Institute of Applied Economic and Social Research
The University of Melbourne
Victoria 3010 Australia
***Telephone* (03) 8344 3701**
***Fax* (03) 8344 5630**
***Email* melb-inst@unimelb.edu.au**
***WWW Address* http://www.melbourneinstitute.com**

Abstract

This paper has three main objectives. First, we re-examine some recent findings that suggest a structural decline in the variance of GDP growth in the United States. We estimate a univariate model in which both the mean growth rate of GDP and its variance are influenced by latent state variables that follow independent Markov chain processes. We are particularly interested in evidence of increased stability in the U.S. economy, either because of reduced volatility or a narrower gap between growth rates in expansions and recessions. Second, we investigate whether a similar phenomenon has occured in other countries. Finally, we explore the extent to which this more general model is better able to describe the shape of actual business cycles.

We find evidence of a reduction in GDP volatility in U.S. data, beginning in late 1984. However, it is less clear that this change represents a structural break. The recent U.S. recession has reduced the probability of being in the low-variance state. Using data from Australia, Canada, Germany, Japan and the United Kingdom, we find evidence of a similar reduction in volatility of GDP growth. The shift for Japan apparently happened in about 1974, and the past decade's poor economic performance seems to have brought a return to the high-variance state. Apart from Germany, the variance reductions in the other countries all occurred within a ten year period between the early 1980's and the early 1990's.

Finally, when we test for non-linear effects using Bayes factors, we find that allowing for a switching variance is much more important than a switching mean. Although the hypothesis of homoscedasticity is overwhelmingly rejected, there is little evidence that this model is better able to capture the shape of actual business cycles.

JEL classification: E32, E37, C22

Keywords: Business cycles, volatility, Markov switching, Bayes factor

1. Introduction and motivation

A decline in the volatility of the output in many of the world's major economies since the mid-1980's has been considered by many as a triumph of central bankers over the business cycle. Whether or not any such stabilization is a result of nature, in the form of good luck or structural change, or nurture by central bankers is an intriguing question. However, the first step in answering such questions is to document the empirical structure of any apparent stabilization and this paper focuses on the measurement of the reduction in output volatility. In particular, we ask whether the mollification of output growth may be attributed to a decline in the volatility of output, a decline in the difference between mean growth rates in recessions and expansions, or both.

Similar questions have recently been posed by (McConnell and Perez-Quiros 2000), (Kim and Nelson 1999a) and (Blanchard and Simon 2001) for the United States and (Mills and Wang 2000) for the G7 countries. Each of these papers models output growth as a univariate autoregressive process. Blanchard and Simon estimate a rolling regression over twenty quater periods and investigate the behavior of the standard deviation of the residual, while the other papers employ a Markov switching approach in which both the mean growth rate and residual variance are driven by independent, state variables.

While all authors document strong evidence of a decline in the volatility of output growth during the post-World War II periods there is some disagreement over the nature of this change. (McConnell and Perez-Quiros 2000), (Kim and Nelson 1999a) and Mills and Wang (2000) find evidence of a one off break in output growth, while Blanchard and Simon (2001) argue that the decline in output volatility has been a gradual process. In addition there are questions about whether there has been a reduction in the difference between mean growth rates. (McConnell and Perez-Quiros 2000) cannot reject the hypothesis that mean growth rates in recessions and expansions are constant across a one off break in the variance of US real GDP growth. However, (Kim and Nelson 1999a) find evidence that there has been narrowing of the gap between mean growth rates in expansions and recessions associated with the decline in the variance of US output growth. The other important difference between the Markov switching models estimated by McConnell and Perez-Quiros and Kim and Nelson is that McConnell and Perez-Quiros allow for general Markov switching in the variance, while Kim and Nelson allow for only a one-time switch in the variance of GDP growth.[1] These differences lead us to ask the following questions. First, since the McConnell and Perez-Quiros model nests that of Kim and Nelson, what is the evidence in favour (or against) this restriction. Second, is there any evidence for similar changes in the behaviour of GDP growth in other countries?

Of this group of papers, only (McConnell and Perez-Quiros 2000) and (Blanchard and Simon 2001) focus on what may have been the cause of the decline in output volatility. McConnell and Perez-Quiros argue that it can be traced to a reduction in the volatility of durable goods output, and in particular to a drop in the share of durable goods accounted for by inventory investment. Blanchard and Simon, however, conclude that there are many "proximate causes" for the stabilisation of output volatility, but identify the more important to be the volatility of inflation and a decrease in consumption and investment volatility.

[1]The other obvious difference between the three Markov Switching papers cited is that McConnell and Perez-Quiros and Mills and Wang use maximum likelihood techniques, in which inferences on the unobserved states are conditional on the point estimates of the other parameters. On the other hand, Kim and Nelson employ Bayesian methods that allow for inference on the states and parameters to be conducted in a symmetric way.

A second major motivation for this paper stems from the work of (Harding and Pagan 2002) and (Hess and Iwata 1997). The ability to produce plausible business cycle features is an important test of any model that purports to explain the business cycle. Both of these papers evaluate several popular models of real GDP growth and find that they generally do not match cyclical characteristics of the observed data. Harding and Pagan develop a new set of nonparametric tools for analysing business cycle characteristics and use them to assess the fit of various models of the cycle, including Hamilton's basic Markov switching model. Amongst their findings is that Markov switching models perform quite poorly relative to a simple AR (1) model. A consequence of this is that Markov-switching non-linear effects do not appear to be very important for describing actual business cycles, despite their popularity and intuitive appeal. To investigate this issue further, we treat the Harding-Pagan statistics as additional functions of interest by mapping the posterior distributions of the model's parameters into posterior distributions of these statistics, based on a simulated data series for each posterior draw. We also exploit a major advantage of our Bayesian estimation methods over classical maximum likelihood: we are able to test for the presence of non-linearity directly through the use of Bayes factors.

To summarise our main findings, we find evidence of a reduction in GDP volatility in U.S. data, beginning in the third quarter of 1984. This is similar to the findings of (McConnell and Perez-Quiros 2000) and (Kim and Nelson 1999a), although they date the volatility reduction from the first quarter of 1984. However, we also find evidence that this is a temporary switch in regime rather than a structural break; the recent U.S. recession has reduced the probability of being in the low-variance state.

Using data from Australia, Canada, Germany, Japan and the United Kingdom, we find evidence of a similar reduction in volatility of GDP growth. The shift for Japan happened in about 1974, but the past decade's poor economic performance seems to have brought a return to the high-variance state. Apart from Germany, the variance reductions in the other countries all occurred within a ten year period between the early 1980's and the early 1990's. The German data clearly shows the effects of reunification in 1991; the combination of West German and unified German data makes the interpretation of this data somewhat problematic.

Finally, when we test for non-linear effects using Bayes factors, we find that allowing for a switching variance is much more important than a switching mean. In a different context, (Sims 2001) and (Sims and Zha 2002) also argue that time-varying volatility produces greater improvements in fit (relative to a linear model) than does time variation in the mean or other coefficients. According to our estimated Bayes factors, the hypothesis that a linear mean process is sufficient to describe GDP growth is roughly an even-money bet, but the hypothesis of homoscedasticity is overwhelmingly rejected. We also use the non-parametric measures of business cycle features recently developed by (Harding and Pagan 2002) to assess the model's fit to the data. Despite the statistical evidence favouring the switching variance model, there is little to suggest that this model is better able to capture the shape of actual business cycles.

2. The model

Our basic Markov switching model is the same as that of (McConnell and Perez-Quiros 2000). The growth rate of real GDP, y_t, follows an autoregressive (AR) process with a switching mean:

$$\phi(L)(y_t - \mu(S_t, D_t)) = e_t, \tag{2.1}$$
$$e_t \sim iidN(0, \sigma^2(D_t)).$$

Here y_t is the first difference of the log of real GDP and $\mu(S_t, D_t)$ is the mean of y_t conditional on the unobserved state vectors S_t and D_t. In addition, the residual variance depends on the value of D_t. Thus the mean growth rate $\mu(S_t, D_t)$ can be affected by the latent process underlying volatility. Specifically, we have

$$\mu(S_t, D_t) = \mu_0 + \mu_{00}D_t + (\mu_1 + \mu_{11}D_t)\, S_t \tag{2.2}$$

and

$$\begin{aligned}\sigma^2(D_t) &= \sigma_0^2(1 - D_t) + \sigma_1^2 D_t \\ &= \sigma_0^2(1 + h_1 D_t),\end{aligned} \tag{2.3}$$

with $h_1 = \left(\frac{\sigma_1^2}{\sigma_0^2} - 1\right)$ .We identify the low-growth state with the event $S_t = 1$ by restricting the mean growth rates in this state, $(\mu_0 + \mu_1)$ and $(\mu_0 + \mu_1 + \mu_{00} + \mu_{11})$, to be negative. This restriction differs from that of (Kim and Nelson 1999a), who imposed the restriction that the mean growth rate in expansions was lower (and that in recessions higher) in the low-variance state than in the high-variance state. In other words, Kim and Nelson impose a narrower gap between the average growth rate in expansions vis-a-vis contractions after the structural break. In addition, we restrict $h_1 < 0$, identifying the low-variance state as $D_t = 1$.

We assume that the latent state variables S_t and D_t are generated by independent first-order hidden Markov chains with transition probabilities $\Pr[S_t = 1|S_{t-1} = 1] = p_{11}$, $\Pr[S_t = 0|S_{t-1} = 0] = p_{00}$, $\Pr[D_t = 1|D_{t-1} = 1] = q_{11}$, and $\Pr[D_t = 0|D_{t-1} = 0] = q_{00}$.The model of (Kim and Nelson 1999a) results from setting $q_{11} = 1$, so that $D_t = 1$ is an absorbing state. Finally, following (Kim and Nelson 1999a) and (McConnell and Perez-Quiros 2000), we specify a first-order autoregression for deviations around the Markov trend, so that $\phi(L) = 1 - \phi L$.

Equation (2.1) is usually estimated by maximum likelihood, however a drawback of this approach is that it requires a degree of approximation when making inferences about S_t and D_t. To see this note that as the state variables are unobserved, estimation of (2.1) is a two stage process. In the first stage, the vector of unknown parameters $\theta = (p_{00}, p_{11}, q_{00}, q_{11}, \phi(L), \mu_0, \mu_1, \mu_{00}, \mu_{11}, \sigma_1^2, \sigma_2^2)'$ is estimated so as to maximize the log of the unconditional density of y_t. This is found to be the sum of the joint distributions across all possible states. When there are only two states under consideration (and conditional on D_t),

$$f(y_t; \theta, D_t) = \sum\nolimits_{j=0}^{1} p(y_t, S_t = j; \theta, D_t), \; j = 0, 1.$$

Once estimates $\hat{\theta}$ of θ have been obtained, inference about the probability of being in a particular state at a given point in time may be made by using the definition of conditional probability:

$$P(S_t = j|y_t; \widehat{\theta}, D_t) = \frac{p(y_t, S_t = j; \widehat{\theta}, D_t)}{f(y_t; \widehat{\theta}, D_t)}.$$

Therefore, estimates of the states do not reflect the uncertainty inherent in the estimates of θ. A Bayesian framework offers an alternative method for making inferences about the state vector. Unlike the classical approach, Bayesian analysis treats both the parameters of the model and the unobserved states as random variables, with inference about S_t drawn from their joint distribution conditional upon the data, $p(S_t, \theta | y_t)$ rather than the conditional distribution, $P(S_t = j | y_t; \widehat{\theta})$.

Recent work by (Albert and Chib 1993)and (McCulloch and Tsay 1994) has demonstrated that Bayesian estimation of Markov switching models is relatively simple to implement using the Gibbs sampler. Gibbs sampling is a Markov chain Monte Carlo (MCMC) method of simulating complex joint and marginal distributions by drawing repeatedly from the conditional distributions, which are much simpler in many cases. As noted by Albert and Chib (1993), the Bayesian approach allows us to treat the unobserved states, $\{S_t, D_t\}_{t=1}^{T}$, as additional parameters to be estimated (through simulation), along with the unknown parameters, θ.

The details of the Gibbs sampling algorithm and the conditional distributions involved are given in (Kim and Nelson 1999a), except that in our model the state variable D_t can be treated in exactly the same way as S_t (i.e., via multi-move rather than single-move sampling; see (Kim and Nelson 1999c)). In our estimation we generate 11,000 iterations and use the final 10,000 for inference. Our prior distributions and starting values are discussed below, in section 3.

In addition to estimating the parameters of the model, we are also interested in testing whether this non-linear model fits the data better than the linear alternative. The model in equations (2.1) to (2.3) reduces to a linear $AR(1)$ if $\mu_1 = \mu_{00} = \mu_{11} = 0$, and $\sigma_1^2 = \sigma_2^2$. As is now well-known, standard likelihood ratio tests cannot be used in this situation because of Davies' problem (the existence of nuisance parameters that are not identified under the null of linearity), and methods such as those of (Hansen 1992) or (Hansen 1996) must be employed. In contrast, Bayesian tests of this hypothesis using Bayes factors are fairly straightforward. As shown in (Koop and Potter 1999), the evidence in favor of linearity can be assessed using the Savage-Dickey Generalized Density Ratio (also see (Verdinelli and Wasserman 1995)). Following (Koop and Potter 1999), we compute Bayes factors in favor of μ_1, μ_{00}, and μ_{11} being equal to zero (individually and jointly), as well as in favor of $\sigma_1^2 = \sigma_2^2$.

Finally, we are also interested in how well this model captures various business cycle features. In particular, we assess the model's performance by computing the non-parametric statistics described in (Harding and Pagan 2002). These are: the average duration and amplitude of expansions and contractions; the average cumulative output gain (loss) during expansions (contractions), and the average 'excess' output gained or lost relative to the 'triangle' approximation to business cycle phases.[2] (Hess and Iwata 1997) conduct a similar exercise, focusing on the amplitude and duration of cycles. Following (Harding and Pagan 2002), we can illustrate their use with the aid of the stylised recession in figure 4.1. In this figure, points A and C represent the peak and trough of the recession, respectively. The duration of the contraction (in quarters) and its amplitude (in percent) are easily interpreted. The cumulated output loss is the area of the triangle ABC plus the area above the actual path of GDP and the line AC. The excess output loss is the difference between the cumulated loss and the triangle area.

We treat the Harding-Pagan statistics as additional functions of interest for our posterior simulator, to use the terminology of (Geweke 1999). Given the output from our posterior simulator, we generate a data series implied by each draw of θ (conditional on

[2]These measures are described in detail in (Harding and Pagan 2002).

the actual initial conditions), and compute the amplitudes, durations, etc. across Monte Carlo draws. Thus, our method incorporates parameter uncertainty into the simulation-based analysis of (Harding and Pagan 2002).

The information we have about the expected duration of business cycles is employed and the prior distributions of p_{00} and p_{11} are accordingly set to have means of 0.8 and standard deviations of 0.16. We use specify relatively non-informative prior distributions with means of 0 and standard deviations of 1 for the other parameters, with the exception of the prior mean for the low growth, high variance state, μ_1, which is set to -0.5.[3]

To start the process of iteration, we set the starting values of p_{00}, p_{11}, q_{00} and q_{11} at 0.9, 0.76, 0.9 and 0.76, respectively. Given these values, we constructed initial state vectors S_t^0 and D_t^0 via the implied Markov process. The starting values for the remaining parameters were computed by the least squares regression of y_t on a constant, its first lag, S_t^0 and D_t^0. A similar procedure for determining starting values is described in Albert and Chib (1993, p. 8). To help ensure that the final estimates were not simply artifacts of the starting values 11,000 draws of the simulation are taken, the first 1,000 of which were discarded.

3. Results

3.1. Parameter estimates

[3] Adrian Pagan has pointed out to us that this prior specification implies a negative average growth rate for GDP. However, as the prior standard deviation for average growth is relatively large we do not feel that this has undluy influenced our results. To check this we re-estimated our models with alternative priors that implied a positive average growth rate and found no substantial differences to our results. We also note that (Kim and Nelson 1999b) use the same prior for de-meaned GDP growth, implying average growth that is below trend.

The results of estimation of (2.1) are presented in table 1, along with our prior means and standard deviations. The Markov switching models tend to imply somewhat longer average phases than we observe in the data. For example, the expected duration of the high-growth state for Australia (given by $(1 - p_{00})^{-1}$) is 30 quarters.[4] The average length of expansion periods in the actual data is 17.5 quarters, based on the BBQ dating algorithm of (Harding and Pagan 2002) (see table 4 below). Similarly, the implied average length of the low-growth state for Australia is 3.6 quarters. Comparable expansion and contraction durations for the U.S. are 28.4 and 4.1 quarters implied by the model versus 21 and 3.2 quarters in the data. For the U.K. the corresponding figures are 32.3 and 4.3 quarters (model) compared with 12.3 and 2.6 quarters (data).

The data are not very informative about q_{00}, the probability of remaining in the high-variance state. The posterior mean estimates are all very close to the prior means of 0.9988 (based on a beta(80,0.1) prior as in (Kim and Nelson 1999a)), while the posterior standard deviations have increased for all countries. Nevertheless, the last two rows of table 1 show a dramatic change in the variance estimates for all six countries. For the United States, our variance estimates are virtually identical to those of (Kim and Nelson 1999a) (see their table 4). The reduction in the variance for the U.S. is not as dramatic as that found by (McConnell and Perez-Quiros 2000); the volatility drops by a factor of 3.7 compared to a more than six-fold decrease in their paper. Relative to the high-state variances σ_0^2, the variance estimates σ_1^2 are lower by a factor of 6.2 for Australia, 4.1 for Canada, 5.6 for Germany 2.5 for Japan, and 6.0 for the UK.

The (smoothed) probabilities of being in the low-variance state are shown in figure 4.2.

[4] Note that this refers to the expected duration of the *model* states, S_t. Some type of dating rule must be used in order to map the S_t into the *business cycle phases* of expansions and recessions. A common choice is that recessions are defined by $\Pr(S_t = 1|y_T) \geq 0.5$.

There is some evidence of clustering in the timing of the variance shifts. Assuming that a switch takes place when $\Pr(D_t = 1|y_T)$ exceeds 50 per cent, the UK, US and Australia all appear to switch into the low-variance state in the early 1980's, although the probability of a shift in the UK remains below 80 per cent until 1992:III. The estimated switching dates for these countries are 1982:II for the UK, 1984:IV for the US, and 1984:III for Australia. The shift in Canada occurs in 1992:I. The variance of Japanese GDP growth enters the low state in 1975:IV, and seems to have switched back to the high state in 1997:III. Japan and Germany are the only countries to display evidence of a switch out of the low-variance state, although the other four (notably the US) display a greater degree of uncertainty (i.e., a larger fall in $\Pr(D_t = 1|y_T)$) near the end of the sample.

Our results regarding the variance shift in the US provides an interesting contrast to the findings of (Blanchard and Simon 2001). First, the estimated probabilities in figure 4.2 are more suggestive of a fairly sharp break in the volatility of GDP, rather than a slow decline. We would expect the latter phenomenon to show up as a more gradual increase in $\Pr(D_t = 1|y_T)$ over the sample. Secondly, when Blanchard and Simon investigate the effects of NBER recessions on their results using a dummy variable to capture these periods, they note that "[o]utput volatility is indeed lower in recessions (by construction)." Our results, on the other hand, suggest that recessions (or more probably, turning points) are associated with higher volatility; the inferences about the variance state are 'clouded' by the recessions of 1991 and 2001. The presence of only one recession between 1983:I and 2000:IV, compared with five between 1960:II and 1982:IV, may well explain part of this result.

3.2. Growth and volatility

The posterior distributions of the mean growth rates in each of the four possible states are summarised in table 4.2. Recall that (Kim and Nelson 1999a) impose the restriction that the gap between the mean growth rates in expansions and contractions is narrower when the variance is low (i.e., after the structural break) than when it is high. In our model, this set of restrictions corresponds to $\mu_{00} < 0$ and $\mu_{00} + \mu_{11} > 0$.[5] Tables 4.1 and 4.2 provide little support for these restrictions. Although the estimates of μ_{00} are negative in Canada, Germany, Japan and the U.S., the standard deviations suggest that much of the posterior distribution lies above zero in each case. According to table 2 the mean growth rate in expansions for the U.S. may have declined in the low-variance period, but not significantly so. The Canadian and Japanese data show this phenomenon much more clearly. The estimated mean growth rate in expansions falls from 1.07 per cent per quarter to 0.77 per cent for Canada, and from over 2 per cent to 0.95 per cent for Japan. The case of Germany is somewhat more ambiguous; although the magnitude of the growth rate decline is the largest in the table, the estimate of μ_0 is very imprecise. This may be due to the effect of reunification.

Except for the low-growth, low-variance state ($S_t = D_t = 1$), our growth rate estimates for the U.S. differ substantially from those of (Kim and Nelson 1999a) (see their table 4). Part of the difference is undoubtedly due to differences in data and sample periods. In particular, Kim and Nelson work with demeaned growth rates whereas we do not.

None of the other countries show any evidence that business cycles have become milder in the sense of a narrowing in the gap between growth rates. On the contrary, mean

[5] Although we use the same notation as (Kim and Nelson 1999a), our parameterization differs from theirs.

growth rates in the high-growth, low-variance state (i.e., when $S_t = 0$ and $D_t = 1$) in Japan and the UK have increased significantly, while Australia shows virtually no change. Interestingly, the average rate of output decline in low-growth periods has increased (the average growth rate has fallen) in all countries. Although this change is not significant, it suggests that recessions may in fact have become more severe, rather than milder. Note that the parameter μ_{11} measures the extent to which the business cycle has become milder, in the sense that the gap between the mean growth rates in expansions and contractions has narrowed. To see this, denote the gaps in the high- and low-variance states as g_0 and g_1, and write their difference as:

$$\begin{aligned} g_0 - g_1 &= [\mu_0 - (\mu_0 + \mu_1)] - [(\mu_0 + \mu_{00}) - (\mu_0 + \mu_{00} + \mu_1 + \mu_{11})] \\ &= -\mu_1 - [-(\mu_1 + \mu_{11})] \\ &= \mu_{11.} \end{aligned}$$

We can therefore base inference about the moderation of business cycles on the posterior distribution of μ_{11}. Apart from the mean and standard deviation of this distribution given in table 4.1, we report the upper-tail area in the bottom row of table 4.2. This is the probability that $\mu_{11} \geq 0$, implying a more moderate business cycle. Evidence for this sort of change is strongest in Germany and Japan, where the probability of a smaller difference in mean growth rates exceeds 68%. On the other hand, the UK business cycle has likely become *less* stable in this sense; roughly 64% of the distribution of μ_{11} lies to the left of zero. For the other countries there is little evidence of any change in the difference in mean growth rates.

3.3. Bayes factors and nonlinearity

Using the Savage-Dickey Density Ratio to compute Bayes factors, we can assess the support in the data for the non-linearities implied by our model. Table 4.3 presents the Bayes factors in favour of the following hypotheses: $\mu_1 = 0$, $\mu_{00} = 0$, $\mu_{11} = 0$, $\mu_1 = \mu_{00} = \mu_{11} = 0$ (labeled 'linear mean' in the table), and $\sigma_0^2 = \sigma_1^2$. Values in excess of one suggest support for the hypothesis in question. (Kass and Raftery 1995) present a rule of thumb for interpreting Bayes factors. Values less than 3 imply that evidence for the hypothesis under study is worth only a bare mention; Bayes factors between 3 and 20 constitute 'positive' evidence; 'strong' evidence corresponds to values between 20 and 150, while Bayes factors in excess of 150 indicate 'very strong' evidence. Using this guideline, there is little evidence for (or against) Markov-switching non-linearity in the mean growth rates for most of these countries. The Bayes factors in table 4.3 suggest that, with the exception of Japan, a linear mean is basically an even-money bet. The evidence *against* linearity in the mean for Japan is 'strong', with a Bayes factor of $1/0.007 = 145$. This result is primarily due to the sharp drop in the mean in the high growth state; the Bayes factor in favor of $\mu_{00} = 0$ is 0.009, or 112 to one against.

The Bayes factors in favor of $\mu_{11} = 0$ are in line with the results discussed above regarding the gap between the mean growth rates.

In marked contrast to the results for the mean parameters, there is very strong evidence against homoscedasticity. Interestingly, it appears that allowing for Markov switching in the residual variance weakens the evidence for Markov switching in the mean of GDP growth.

3.4. The shape of business cycles

Our final set of results concerns the non-parametric statistics of (Harding and Pagan 2002) that describe the shape of the business cycle. By treating these measures as additional functions of interest, we can obtain their posterior distributions in a straightforward way. These distributions, for both recessions and expansions, are presented in figures 4.3 to 4.14. We dated the periods of expansion and recession using the 'BBQ' dating algorithm described in (Harding and Pagan 2002). In each figure, the vertical line indicates the true value computed from the data (i.e., the average of the given statistic over the observed peaks and troughs). The distributions from our base model (with both switching mean and variance) are represented by a solid line and labelled 'b' in the figures. The dashed line ('m') corresponds to a model with a switching mean but constant variance, while the dotted line ('v') arises from a model with a linear mean and switching variance. We investigate this last model based on the (admittedly inclusive) evidence in favour of linearity in the mean given by the Bayes factors discussed above. Finally, we include a model (shown by the dash-dot line, labelled 'i') with both switching mean and variance, but with $\mu_{00} = \mu_{11} = 0$. In this model, the state of the variance has no impact on the state of the mean. This model also receives some support in table 4.3.

Several interesting points emerge from an examination of these figures. First, the ability of simulated data from these models to adequately capture the shape of actual business cycles is generally quite good, with modal values near the average values observed in the data. There are some notable exceptions however. For example, the Markov switching models significantly underestimate all aspects of Australian expansions (see figure 4.4).

Second, the distributions of the Harding-Pagan statistics are often very skewed, with extremely long tails (the x-axes in some of the figures have been truncated to show the central mass of the distributions).[6] Clearly, the modal estimates of the distributions are much more favourable to the Markov switching models than are the means and standard deviations. For example, the posterior distribution of the duration of US expansions, using the base model, has a mean of 26.8 quarters and a standard deviation of 15.5 quarters. One would not reject the hypothesis that the model captured the true value of 16.8 quarters simply because of the large standard deviation. On the other hand, the true value lies at roughly the 23^{rd} percentile of the posterior distribution, while the mode of 19.4 quarters lies at the 33^{rd} percentile. The fit of the model is not great, but is better than the 't-statistic' approach would suggest.

Third, in most cases there is little to choose from between the three models, in the sense that the modes of the various distributions are roughly coincident and/or equally far from the data values. Where there is a noticeable difference, there is no consistent ranking between the three models. For Canadian recessions, the model with only a switching mean is closest to the actual value of the amplitude, duration and cumulative loss. Next is the model with independently switching mean and variance, while the variance-only model is furthest. For Australian recessions the model with independently switching mean and variance does slightly better than the others.

Finally, the posterior distributions of the 'excess' statistics are much more symmetric than those of the other statistics, with modes that are closer to zero than the actual data values. Recall that an 'excess' value of zero implies a linear growth path for GDP, or a perfect fit for the triangle approximation of (Harding and Pagan 2002). The interesting

[6] The maximum truncation occurred in the "cumulative gain" panel of figure 4.8, where 17% of the posterior of the variance-only model lies to the right of 1,000. In most other cases the truncation was less than 2% of the distribution. Further details are available on request.

feature of these posteriors is that the four variations of the Markov switching models all imply virtually the same degree of departure from the triangle approximation, and so are all about equally informative (or not) about this particular version of nonlinearity in the data.

4. Conclusions and future directions

We draw three main conclusions from the results in this paper. First, there seems to have been a general change in the volatility of GDP growth in the countries we study in the mid-1980's. For all countries except Japan, this change was toward decreased volatility. It is not clear, however, that this has been a permanent structural shift.

Second, the evidence is much stronger for non-linearity in the variance of GDP growth than it is for non-linearity in its mean. Bayes factors overwhelmingly reject homoscedasticity, but suggest that a linear mean growth rate is a slightly better than even-money bet.

Third, the models we examine are able to generate data that does a reasonable job of replicating the shape of actual business cycles. The models' ability to match the non-parametric statistics of (Harding and Pagan 2002) varies across countries and phases of the cycle, and becomes most apparent when one examines the entire posterior distribution. In future work, we plan to investigate the possibility of using non-parametric shape measures such as the Harding-Pagan statistics to elicit 'business cycle priors' for regime-switching models, including those presented here.

References

Albert, J. H., and S. Chib (1993), 'Bayes Inference Via Gibbs Sampling of Autoregressive Time Series Subject to Markov Mean and Variance Shifts,' *Journal of Business and Economic Statistics* 11(1), 1–15.

Blanchard, O., and J. Simon (2001), 'The Long and Large Decline in U.S. Output Volatility,' *Brookings Papers on Economic Activity* 2001(1), 135–174.

Geweke, J. (1999), 'Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication,' *Econometric Reviews* 18(1), 1–73.

Hansen, B. E. (1992), 'The Likelihood Ratio Test under Non-Standard Conditions: Testing the Markov Switching Model of GNP,' *Journal of Applied Econometrics* 7, S61–S82.

——— (1996), 'Inference When a Nuisance Parameter is Not Identified under the Null,' *Econometrica* 64(2), 357–384.

Harding, D., and A. R. Pagan (2002), 'Dissecting the Cycle: A Methodological Investigation,' *Journal of Monetary Economics* 49(2), 365–381.

Hess, G. D., and S. Iwata (1997), 'Measuring and Comparing Business-Cycle Features,' *Journal of Business and Economic Statistics* 15(4), 432–444.

Kass, R. E., and A. E. Raftery (1995), 'Bayes Factors,' *Journal of the American Statistical Association* 90(430), 773–795.

Kim, C.-J., and C. R. Nelson (1999a), 'Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of Business Cycle,' *Review of Economics and Statistics* 81(4), 608–616.

——— (1999b), 'Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of Business Cycle,' *Review of Economics and Statistics* 81(4), 608–616.

——— (1999c), *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, MIT Press, Cambridge, MA.

Koop, G., and S. M. Potter (1999), 'Bayes Factors and Nonlinearity: Evidence from Economic Time Series,' *Journal of Econometrics* 88, 251–281.

McConnell, M. M., and G. Perez-Quiros (2000), 'Output Fluctuations in the United States: What Has Changed Since the Early 1980's?,' *American Economic Review* 90(5), 1464–1476.

McCulloch, R. E., and R. S. Tsay (1994), 'Statistical Analysis of Economic Time Series Via Markov Switching Models,' *Journal of Time Series Analysis* 15(5), 523–538.

Mills, T. C., and P. Wang (2000), 'Searching for the Sources of Stabilisation in Output Growth Rates: Evidence from the G-7 Economies,' *Research paper no. 00/7, Department of Economics, Loughborough University* .

Sims, C. A. (2001), 'Stability and Instability in US Monetary Policy Behavior,' *mimeo, Princeton University, www.princeton.edu/ sims* .

Sims, C. A., and T. Zha (2002), 'Macroeconomic Switching,' *mimeo, Princeton University, www.princeton.edu/ sims* .

Verdinelli, I., and L. Wasserman (1995), 'Computing Bayes Factors Using a Generalization of the Savage-Dickey Density Ratio,' *Journal of the American Statistical Association* 90, 614–618.

Table 4.1: Prior and posterior distributions of model parameters

		Prior	Australia	Canada	Germany	Japan	UK	US
p_{00}	mean	0.8	0.9667	0.9742	0.9506	0.9677	0.9690	0.9648
	s.d.	0.16	0.0269	0.0181	0.0518	0.0197	0.0227	0.0234
p_{11}	mean	0.8	0.7211	0.7762	0.6678	0.8827	0.7695	0.7547
	s.d.	0.16	0.1497	0.1314	0.1748	0.1071	0.1353	0.1343
q_{00}	mean	0.9988	0.9926	0.9923	0.9722	0.9911	0.9863	0.9910
	s.d.	0.004	0.0072	0.0077	0.0183	0.0084	0.0129	0.0088
q_{11}	mean	0.9	0.9836	0.9712	0.9748	0.9810	0.9733	0.9763
	s.d.	0.091	0.0160	0.0284	0.0156	0.0169	0.0224	0.0237
μ_0	mean	0	0.9451	1.0652	1.0012	2.002	0.5892	0.9349
	s.d.	1	0.1784	0.1466	0.5301	0.2286	0.1700	0.1875
μ_1	mean	-0.5	-1.2461	-1.3346	-1.3793	-2.1488	-0.9132	-1.1734
	s.d.	1	0.3278	0.2807	0.5276	0.2690	0.3249	0.2728
μ_{00}	mean	0	0.0008	-0.2954	-0.3444	-1.0538	0.1378	-0.0880
	s.d.	1	0.2018	0.2045	0.5634	0.2679	0.1880	0.2194
μ_{11}	mean	0	-0.0385	-0.2232	0.2865	0.8183	-0.2333	-0.1604
	s.d.	1	0.5302	0.8287	0.6774	0.7562	0.6080	0.6527
ϕ_1	mean	0	-0.0146	0.2406	-0.0994	-0.0617	0.0317	0.1526
	s.d.	0.5	0.0909	0.1074	0.0932	0.1100	0.0966	0.1256
σ_0^2	mean	0.333	2.8590	0.9782	8.7798	1.6242	1.5738	1.0195
	s.d.	0.408	0.4618	0.1567	4.3577	0.3653	0.4439	0.2185
σ_1^2	mean	0.333	0.4620	0.2362	1.5794	0.6608	0.2623	0.2720
	s.d.	0.408	0.1051	0.0890	0.6217	0.2020	0.0999	0.0798

Table 4.2: Mean growth rates by state

S_t	D_t			Prior	Australia	Canada	Germany	Japan	UK	US
0	0	μ_0	mean	0	0.945	1.065	1.001	2.002	0.589	0.935
			s.d.	1	0.178	0.147	0.530	0.229	0.170	0.188
1	0	$\mu_0+\mu_1$	mean	-0.5	-0.301	-0.269	-0.378	-0.147	-0.324	-0.239
			s.d.	1.41	0.293	0.252	0.368	0.152	0.303	0.249
0	1	$\mu_0+\mu_{00}$	mean	0	0.946	0.770	0.657	0.948	0.727	0.847
			s.d.	1.41	0.098	0.136	0.137	0.134	0.077	0.098
1	1	$\sum \mu$	mean	-0.5	-0.339	-0.788	-0.436	-0.382	-0.420	-0.487
			s.d.	2	0.501	0.816	0.556	0.696	0.561	0.639
			$pr(\mu_{11} \geq 0)$		0.509	0.481	0.689	0.890	0.376	0.493

Table 4.3: Bayes factors

	Hypothesis				
	$\mu_1 = 0$	$\mu_{00} = 0$	$\mu_{11} = 0$	linear mean	$\sigma_0^2 = \sigma_1^2$
Australia	1.3603	5.0603	1.2338	1.9102	3.26E-14
Canada	0.4659	1.3548	1.2831	0.975	6.87E-05
Germany	1.4548	1.2239	1.2042	1.9635	1.23E-25
Japan	0.2475	0.0089	0.6857	0.0069	2.05E-24
UK	1.1469	4.371	1.4846	1.4551	6.98E-17
US	0.6228	3.9828	1.7245	1.8696	6.54E-06

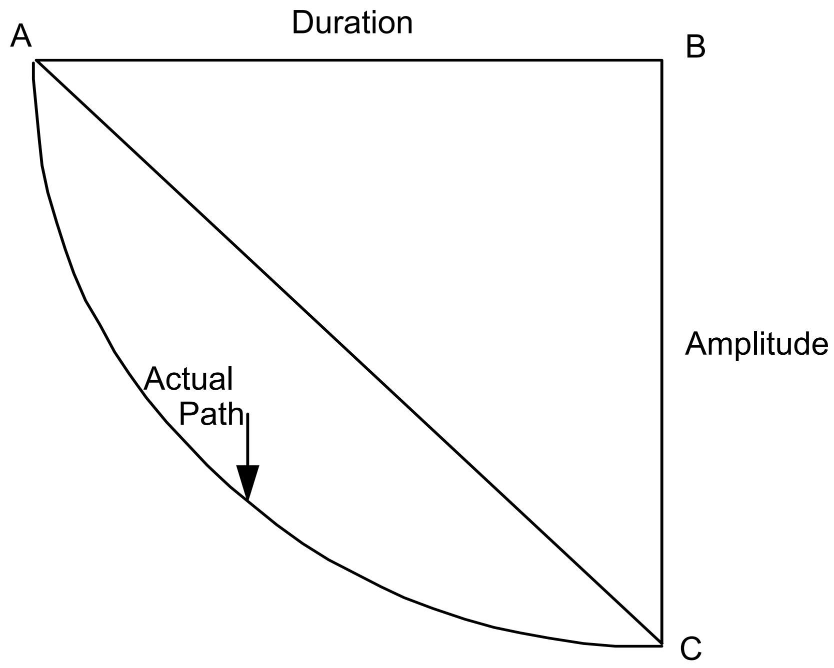

Figure 4.1: Sylised recession phase

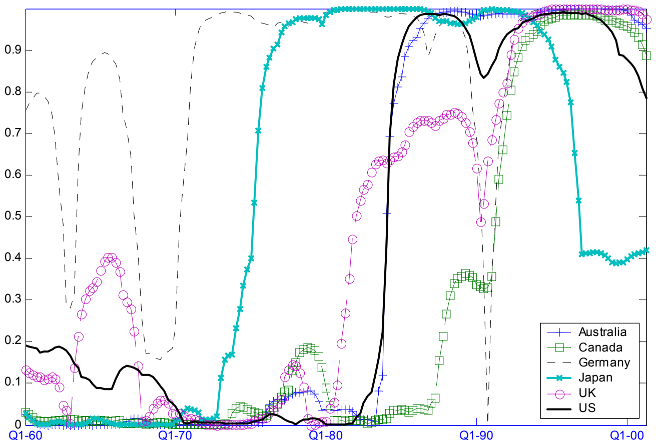

Figure 4.2: Probability of being in low-variance state

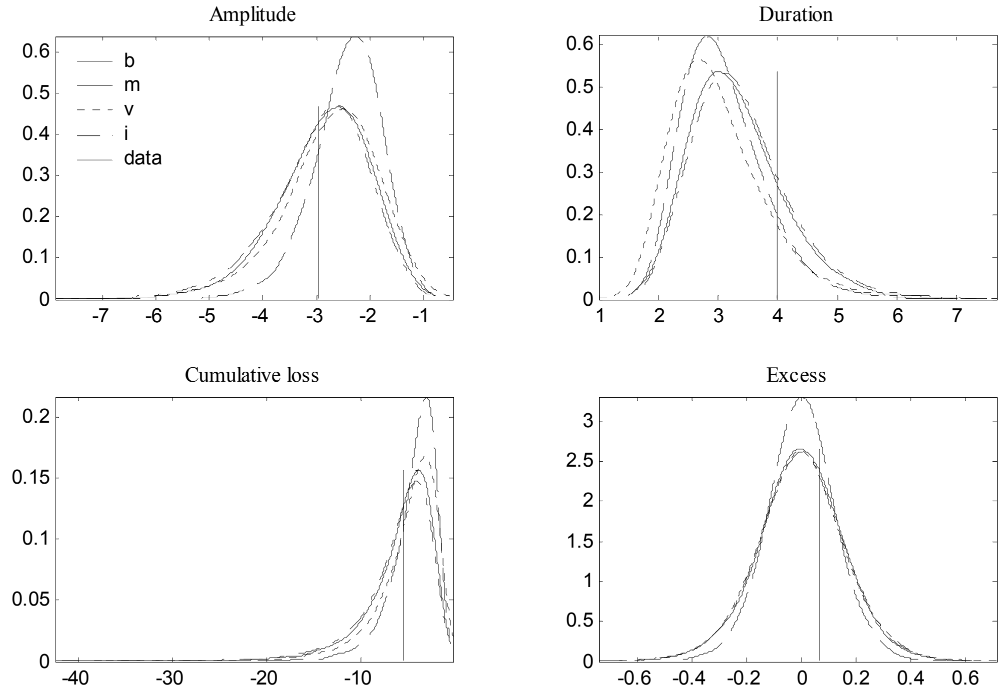

Figure 4.3: Harding-Pagan statistics for Australian recessions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

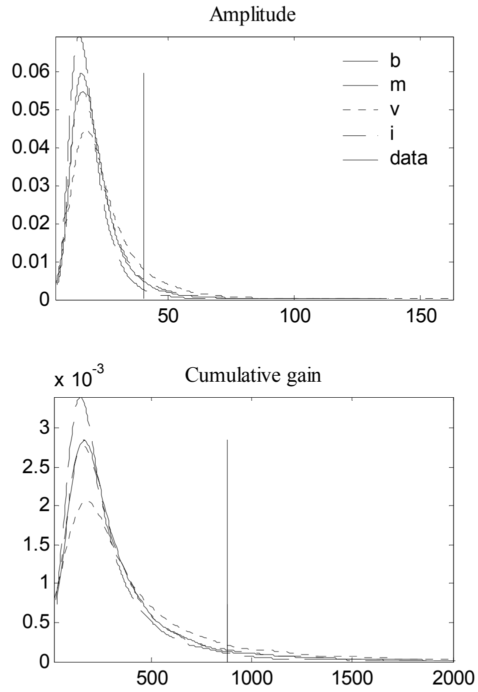

Figure 4.4: Harding-Pagan statistics for Australian expansions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

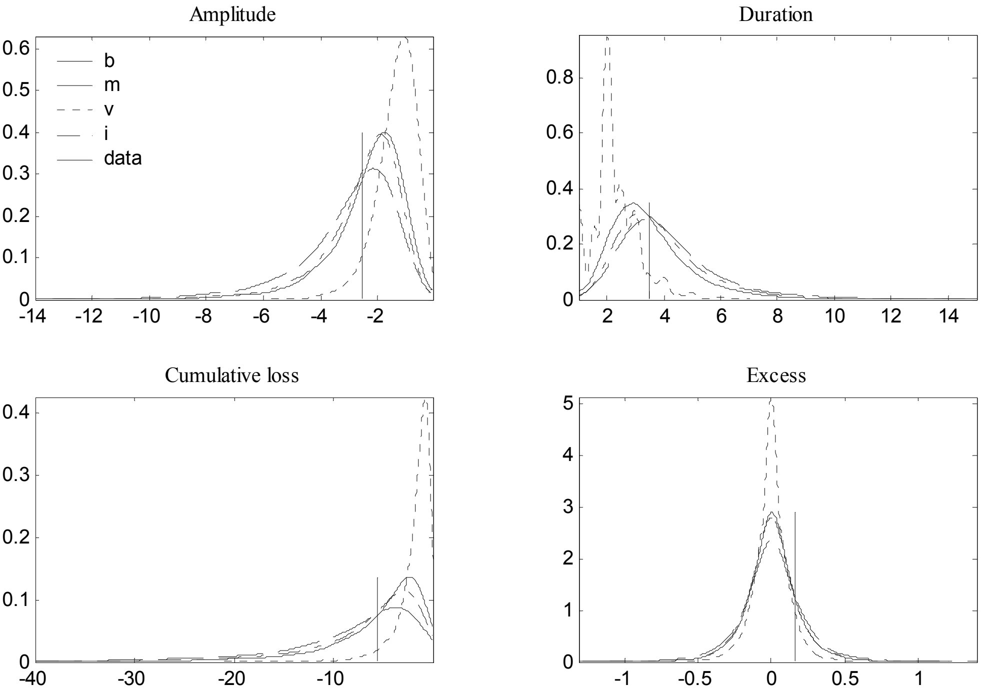

Figure 4.5: Harding-Pagan statistics for Canadian recessions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

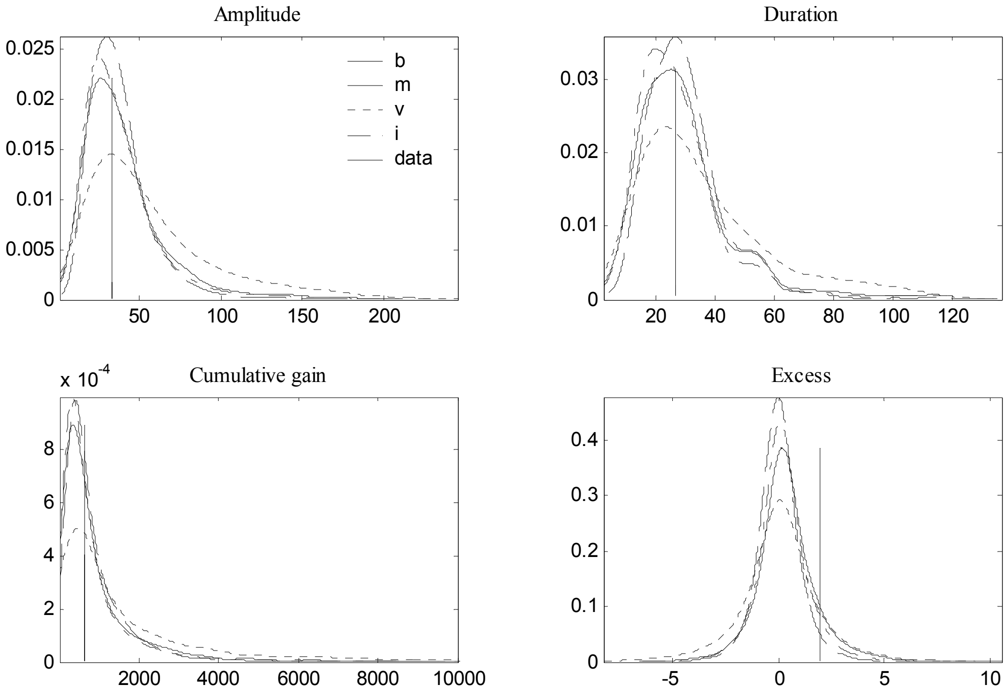

Figure 4.6: Harding-Pagan statistics for Canadian expansions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

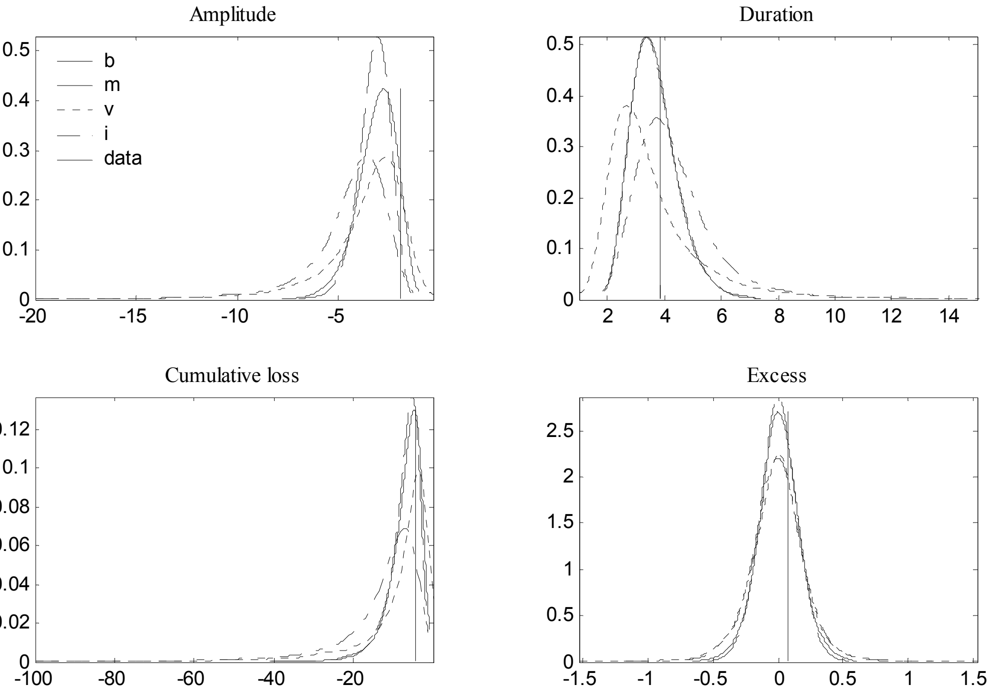

Figure 4.7: Harding-Pagan statistics for German recessions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

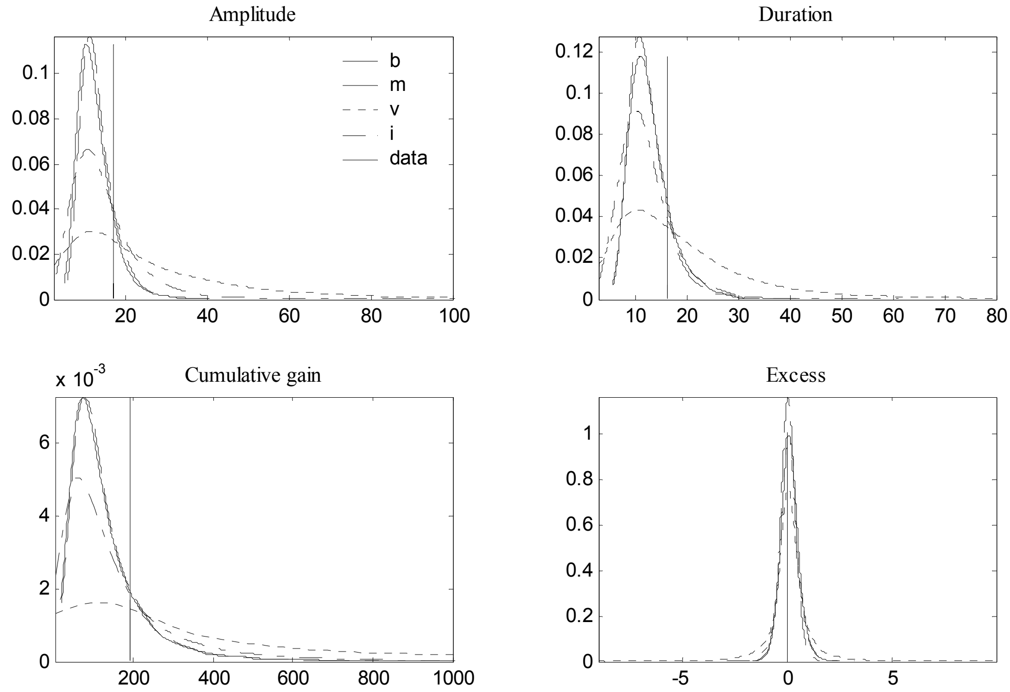

Figure 4.8: Harding-Pagan statistics for German expansions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

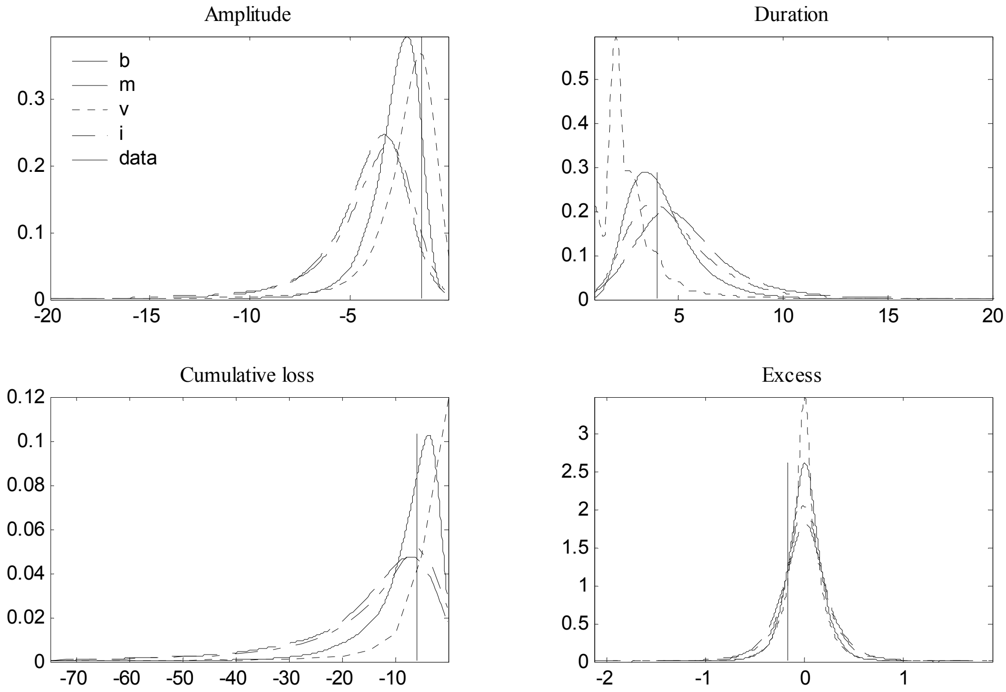

Figure 4.9: Harding-Pagan statistics for Japanese recessions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

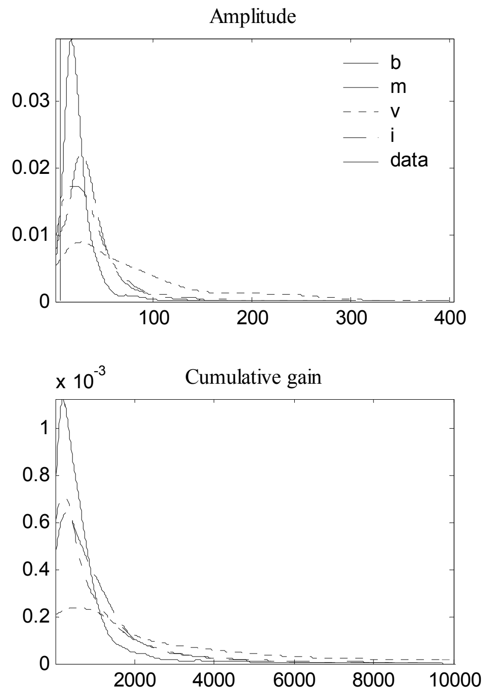

Figure 4.10: Harding-Pagan statistics for Japanese expansions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

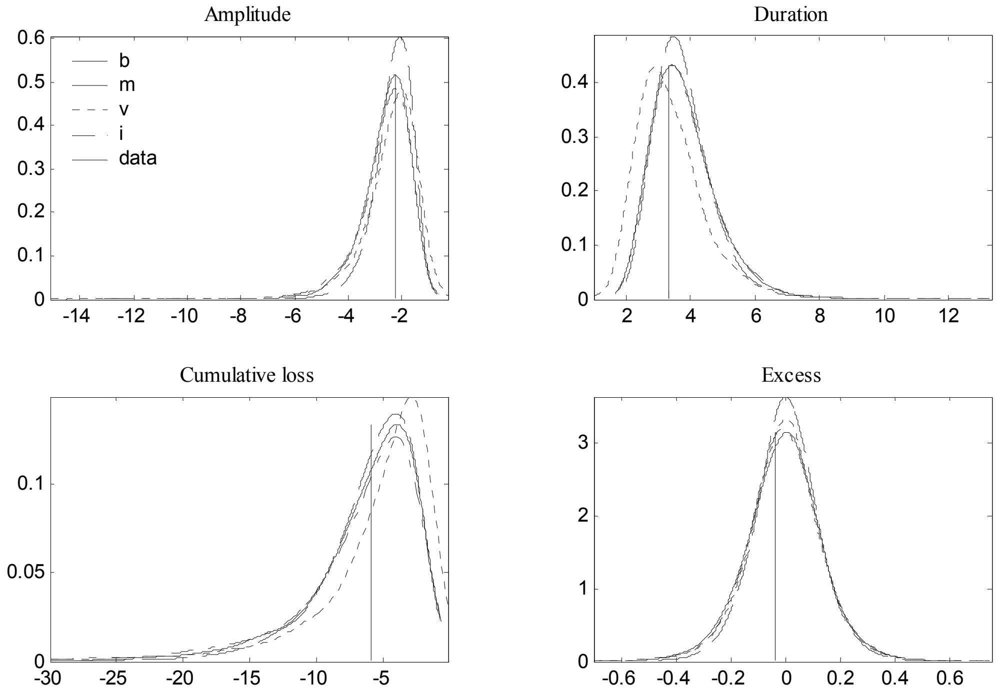

Figure 4.11: Harding-Pagan statistics for United Kingdom recessions. Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

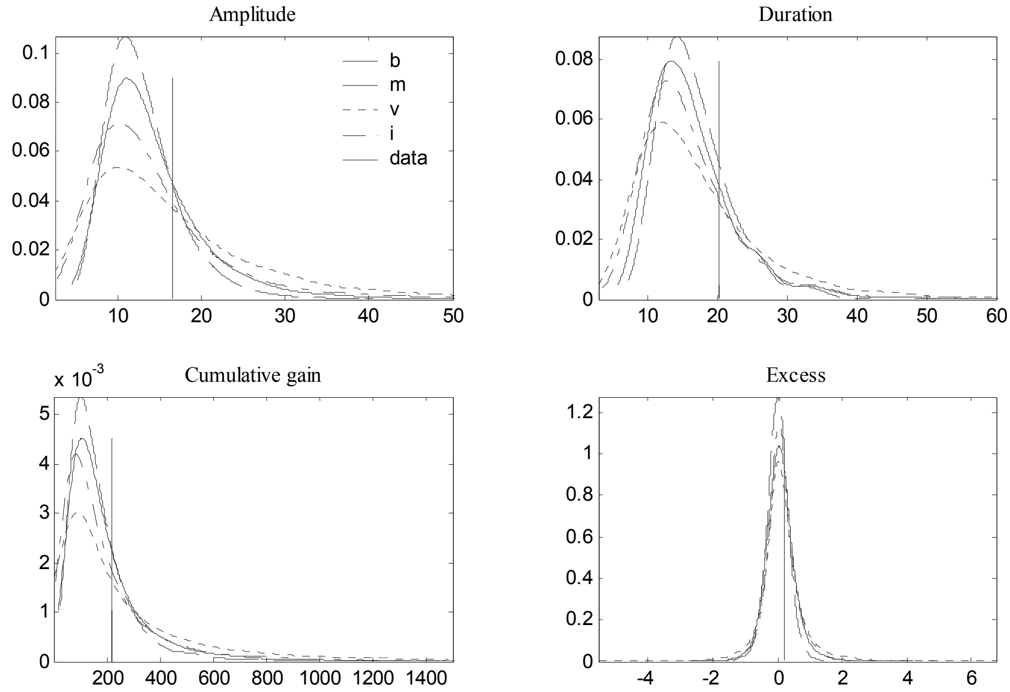

Figure 4.12: Harding-Pagan statistics for United Kingdom expansions. Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

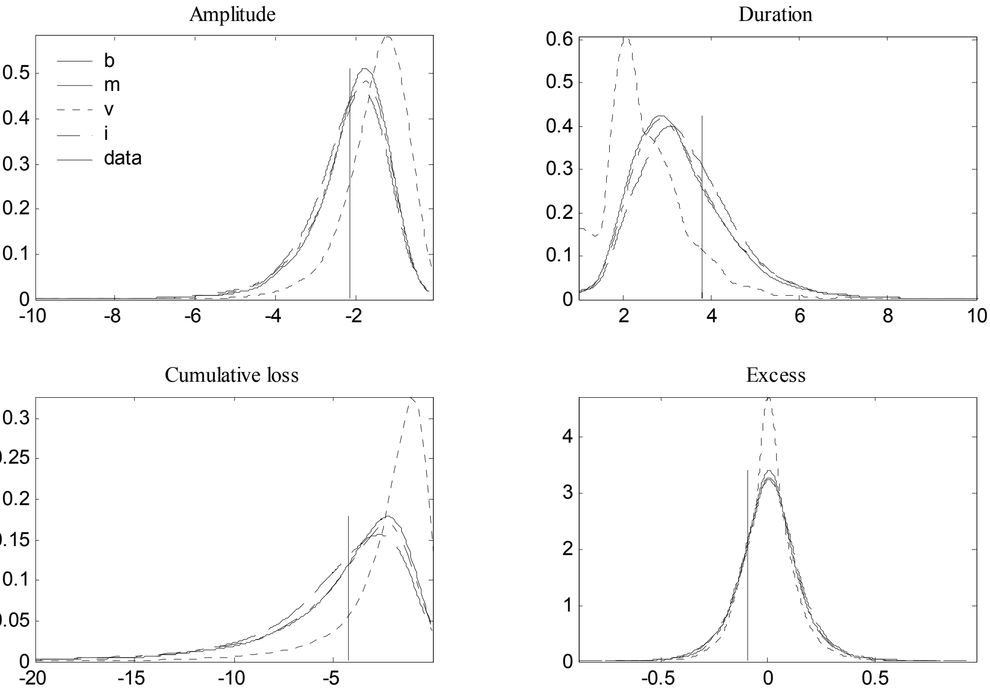

Figure 4.13: Harding-Pagan statistics for United States recessions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.

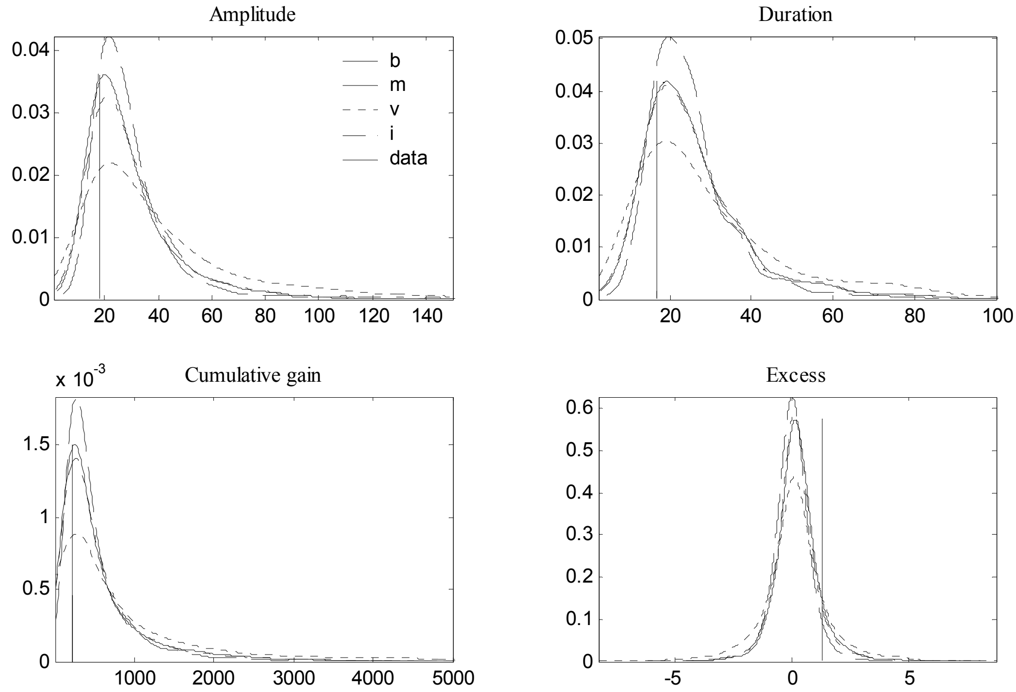

Figure 4.14: Harding-Pagan statistics for United States expansions.
Models : b – switching mean and variance; m – switching mean only; v – switching variance only; i – model b with mean and variance independent.