Studies in Nonlinear Dynamics and Econometrics

Quarterly Journal
April 1996, Volume 1, Number 1
The MIT Press

Studies in Nonlinear Dynamics and Econometrics (ISSN 1081-1826) is a quarterly journal published electronically on the Internet by The MIT Press, Cambridge, Massachusetts, 02142. Subscriptions and address changes should be addressed to MIT Press Journals, 55 Hayward Street, Cambridge, MA 02142; (617)253-2889; e-mail: journals-orders@mit.edu. Subscription rates are: Individuals $40.00, Institutions $130.00. Canadians add additional 7% GST. Prices subject to change without notice.

Permission to photocopy articles for internal or personal use, or the internal or personal use of specific clients, is granted by the copyright owner for users registered with the Copyright Clearance Center (CCC) Transactional Reporting Service, provided that the per-copy fee of $10.00 per article is paid directly to the CCC, 222 Rosewood Drive, Danvers, MA 01923. The fee code for users of the Transactional Reporting Service is 0747-9360/96 $10.00. For those organizations that have been granted a photocopy license with CCC, a separate system of payment has been arranged. Address all other inquiries to the Subsidiary Rights Manager, MIT Press Journals, 55 Hayward Street, Cambridge, MA 02142; (617)253-2864; e-mail: journals-rights@mit.edu.

© 1996 by the Massachusetts Institute of Technology

A Check on the Robustness of Hamilton's Markov Switching Model Approach to the Economic Analysis of the Business Cycle

Michael D. Boldin*
The Conference Board
New York, NY
boldin@conference-board.org

Abstract. *This note explores the robustness of Hamilton's (*Econometrica*, 1989) two-regime Markov switching model framework for capturing business-cycle patterns. Applying his exact specification to a revised version of real GNP, I find parameter estimates that are similar to those he reported only when I use the same sample period (1952–1984) and a particular set of starting values for the maximum likelihood procedure. Two other local maxima exist that have higher likelihood values, and neither correspond to the conventional recession-expansion dichotomy. In fact, when the sample period is extended, there is no longer a local maximum near the parameter set reported by Hamilton. Exploring the model and data further, I reject cross-regime restrictions of Hamilton specification, but also find that relaxing these restrictions increases the number of local maxima. However, a parsimonious three-regime model for GNP growth is more robust and plausible, especially when each regime is required to last more than one quarter.*

1 Introduction

James Hamilton's article "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle" (1989) has spawned a considerable number of papers that also use Markov switching models (MSMs) to characterize and explain business-cycle fluctuations. These empirical studies are primarily motivated by a belief that recessions and expansions are distinct phases or regimes that make economic fluctuations a fundamentally asymmetric phenomenon. Because the estimated parameters of relatively simple MSM specifications match many stylized facts about the business cycle, this framework has become an important alternative to linear, autoregressive structures.

This article explores the robustness of maximum-likelihood estimates of Hamilton's MSM specification. I first apply Hamilton's MSM approach to a revised version of his original data series, the growth rate of real GNP. Over the 1952–1984 sample period, I obtain almost identical coefficients for one set of parameters that correspond to a local maximum of the likelihood function. However, two other local maxima exist that have higher likelihood values and do not correspond as closely to the conventional recession-expansion dichotomy (i.e., those based on National Bureau of Economic Research turning points). In fact, when the sample period is extended, there is no longer a local maximum near the parameter set reported by Hamilton.

One problem with Hamilton's MSM specification is that certain cross-regime restrictions are easily rejected by the data.[1] Generalizations of the model tend to increase the number of local maximums in the likelihood

*This paper was completed while I was an economist at the Federal Reserve Bank of New York. I am grateful to Jim Hamilton for providing details about his model. Also, many colleagues in the Federal Reserve System contributed helpful suggestions for presenting this research, but I am responsible for all errors or omissions.

[1]This finding is noteworthy since Lam (1990) generalized the model to allow for a form of trend stationarity in the log-level of GNP, Goodwin (1993) applied Hamilton's exact specification to real GNP growth of seven other OECD countries, and Filardo (1994) only changed the transition process to fit monthly U.S. IP growth. None of these studies considered the importance of the restrictions that I relax.

© 1996 by the Massachusetts Institute of Technology

function, however. This problem arises principally from overparameterization that tends to cause estimates of unrestricted MSMs to stray from plausible business-cycle parameters. I then show that a three-regime model for GNP growth (recession, rapid growth, and normal expansion) is more robust and plausible, especially when each regime is required to last more than one quarter.

2 Hamilton's Model and Maximum Likelihood Estimation

The MSM that Hamilton estimated has two unobservable regimes (indexed by s_t) for GNP growth:

$$y_t = \alpha_1 + u_t \qquad \text{if } s_t = 1$$
$$y_t = \alpha_2 + u_t \qquad \text{if } s_t = 2$$

To identify the regimes, $\alpha_1 < \alpha_2$ is assumed; to match basic recession and expansion characteristics, $\alpha_1 < 0 < \Sigma_t y_t/T < \alpha_2$ is expected.[2] In addition, Hamilton's specification allows for fourth-order autoregressive (AR) effects:

$$u_t = \rho_1 u_{t-1} + \rho_2 u_{t-2} + \rho_3 u_{t-3} + \rho_4 u_{t-4} + e_t$$

to capture either inertia or trend-reversion forces that are the same in both regimes. Also, the distribution of the random variable, e_t, is assumed to normal with the square root of the variance, σ, the same in both regimes.

In an MSM specification, the regime sequence is assumed to follow a random switching process that is independent of e_t and is first-order or Markov where:

$$q_{11} = \text{prob}(s_t = 1 \mid s_{t-1} = 1), \qquad q_{22} = \text{prob}(s_t = 2 \mid s_{t-1} = 2).$$

Note that these two probabilities define the entire Markov switching process, since

$$q_{12} = 1 - q_{11} = \text{prob}(s_t = 2 \mid s_{t-1} = 1), \qquad q_{21} = 1 - q_{22} = \text{prob}(s_t = 1 \mid s_{t-1} = 2).$$

Most important, this setup assumes that recessions and expansions are not duration dependent. Once a new regime begins, the odds of it ending do not vary over time. We would expect persistence in the regime generating process with $q_{ii} > 0.50$.

Because the regimes are unobservable, their effects and incidence must be inferred from the data. Maximizing the likelihood function that accounts for all possible regime sequences is the most direct manner to accomplish this task. In the case where there are no serial correlation effects, the likelihood function is:

$$L(y; \alpha, \sigma, \rho, q) = \sum_{s_T=1}^{2} \sum_{s_{T-1}=1}^{2} \cdots \sum_{s_1=1}^{2} \sum_{s_0=1}^{2} \left[\prod_{t=1}^{T} f_{s,t} q_{s_t, s_{t-1}} \right] p_{s_0}$$

where the α, σ, and ρ parameters define probability density functions (pdfs): $f_{s,t} = f(y_t - \alpha_s, \sigma)$ for each regime at each time period. The likelihood function then weights the pdfs by the appropriate transition parameters, $q_{s(t-1),s(t)}$, and the initial period's regime probabilities, p_{s_0}. In the absence of pre-sample information, the latter parameters are set at the unconditional regime probabilities, which depend on the transition probabilities and also represent the proportion of the observations that are in each regime.[3]

Maximum likelihood (ML) estimation entails using a numerical optimization routine that starts with an initial guess for the MSM parameters and then searches for the set that corresponds to a distinct peak in the likelihood function.[4] But because the likelihood function contains terms for every possible regime sequence,

[2] Hamilton set up the model as $y_t = \alpha_0 + \alpha_1 s_t + u_t$ with $s_t = 0$ or 1, which is equivalent to the framework that I use but is more difficult to generalize to models with more than two regimes.

[3] In the two regime case, the unconditional probabilities are $p_1 = (1 - q_2)/(2 - q_1 - q_2)$ and $p_2 = (1 - q_1)/(2 - q_1 - q_2)$. In general, they are the solution to $p_0 = p_0^* Q$, where Q is the matrix of transition probabilities: $q_{ij} = \text{prob}(s_t = j \mid s_{t-1} = i)$.

[4] In practice, the negative of the log-likelihood is usually minimized and constraints must be placed on the σ and q_{ii} parameters to keep them within the permissible set. For the estimates presented below, I use the GAUSS (PC-DOS) software package with the OPTMUM application program that employs a modified Newton-Raphson method. A set of GAUSS programs for replicating my results and estimating other interesting MSM specifications is available upon request.

which is of the order of 2^T, computing the value for any set of parameters may seem infeasible. However, Cosslett and Lee (1984) showed that the likelihood value can be derived using an algorithm that keeps the calculations to a much smaller order (n^2T, with n being the number of regimes and T being the number of observations).[5]

While it is not necessary to examine all details of the estimation of these models, it is noteworthy that the first-order conditions for $\alpha 1$ and $\alpha 2$, which help define a peak in the likelihood function, satisfy the following weighted-mean equation:

$$\alpha_i = (\Sigma_t y_t p^*_{i,t})/(\Sigma_t p^*_{i,t}),$$

where

$$p^*_{i,t} = \text{prob}(s_t = i \mid y) = L(y, s_t = i; \alpha, \sigma, \rho, q)/L(y; \alpha, \sigma, \rho, q).$$

The latter formula uses Bayes's rule to compute regime probabilities, with the denominator being the model's likelihood value and the numerator being a modification that only considers regime sequences that include $s_t = i$. Most important, both values use all observations in the sample and therefore the p^*s are full-sample regime probabilities.[6]

The first-order conditions for the other parameters can be arranged into analogous formulas, with the full-sample probabilities again serving as intuitively appealing observation weights. Therefore, ML estimates are completely consistent with the p^*s, which are calculated using the ML estimates. The relationships between the ML estimates and the full-sample probabilities for each regime mean that the results most completely and consistently represent how the data fits into a particular multiregime view. This property contrasts with alternative regime analysis techniques that either fix the time periods to estimate the regime characteristics, or instead, fix the regime characteristics to determine the observations that belong to each regime.

3 Evaluating MSM Estimates

Despite the attractiveness of ML estimates, the shape of the likelihood function makes testing MSM results difficult, if not impossible. Multiple peaks and cases where the Hessian (matrix of second derivatives) of the likelihood function is singular means that test distributions that assume an approximately quadratic shape are not very useful. In Hamilton's model, the singularity problem is seen along the line $\alpha_1 = \alpha_2$. This case is of particular interest, since it yields a single-regime model. Unfortunately, the restriction $\alpha_1 = \alpha_2$ cannot be tested using conventional likelihood ratio (LR) and Wald statistics because the distribution of these statistics assumes that a singularity problem does not exist at the null hypothesis. It can also be shown that Lagrange multiplier tests have no power in this case. Furthermore, there is a nuisance parameter problem, because the transition probabilities are irrelevant or unidentified when there is a single regime. Although LR and Wald tests can be modified to control for nuisance parameters, the singularity condition in the information matrix creates an obstacle that cannot be addressed with classical testing procedures.[7]

Another condition that affects the shape of the likelihood function and is relevant for this study pertains to MSM specifications that allow for the variance of e_t to differ across regimes. In this case, the likelihood function becomes unbounded when one regime is identified with only one observation. For say, $s = 2$ and $t = t'$:

$$\text{as} \qquad \alpha_2 \rightarrow y_{t'} \qquad \text{and} \qquad \sigma_2 \rightarrow 0, \qquad f(y_{t'} - \alpha_2, \sigma_2) \rightarrow \infty.$$

[5]Boldin (1992) discusses this algorithm in detail, shows the equivalence to the computationally more intensive, but somewhat more intuitive algorithm that Hamilton used, and describes how to estimate models with autoregressive terms and higher-order transition process by expanding the number of effective regimes to represent a sequence of the original regimes.

[6]Hamilton (1989) refers to these values as "smoother" probabilities and they contrast with a residual output from his likelihood function algorithm that only uses information up to data "t" to define $p_{i,t} = \text{prob}(s_t = i \mid y_1, \ldots y_t)$.

[7]Similar singularity conditions and nuisance parameter problems occur when testing two versus three regimes, or whenever the null hypothesis has less regimes than the alternative.

Although these cases can be ruled out as uninteresting or implausible and the analysis can concentrate on interior ML solutions (local peaks), multiple local maxima are likely and extremely problematic.[8] In practice, this means that it is important to try various starting values in the optimization routine and consider whether each converged-to set of ML parameters is plausible.

These problems mean that simple and direct statistical criteria cannot be used to either accept or reject these models. A more pragmatic approach is needed to judge their contribution to the empirical analysis of business cycles.[9] In this study, I estimate the models over various sample periods and look for reasonable stability in the most interesting parameters. In Hamilton's model, the focus is on α_1 and α_2, which measure the mean GNP growth rate in recessions and expansions, respectively. Also, the transition probabilities are interesting, since they measure the persistence of each stage of the business cycle. Besides looking for sample stability, I consider how both the log-likelihood measured fit and the other parameters change when certain aspects of the model are modified. When these specification changes are made, the improvement in the likelihood function, as measured by the LR statistic, is only used as a guide to decide whether this is an important or significant change, however. Problems of instability as the sample period expands, or finding that a particular regime's σ value is close to zero, gets precedence over any type of LR result.

Finally, I look at the full-sample probabilities for recession or low-growth regimes, and make comparisons to NBER business-cycle dates to determine whether a plausible business-cycle pattern is found. A formal test is not used, however, since this would imply that the NBER dates are the truth, and if this was the case, we should simply use these dates to split the data and estimate the parameters for each regime. Also, it must be recognized that the full-sample probabilities are not an independent source of information since they are strongly tied to the estimated MSM parameters. They do provide a useful summary of the model's distinction between recession and expansion regimes, however.

Admittedly, I am performing a specification search. Since the goal is to find plausible business-cycle patterns from a robust MSM, it should not be surprising that I can report some success in this endeavor. However, I also attempt to simplify the preferred MSM specification. This exercise can be described as a quest to find the simplest MSM specification that both captures the most interesting characteristic of business cycles and is robust to changes in the sample period. An MSM for GNP growth that satisfies this criteria is a good base model that can be expanded upon to study different aspects of the business cycle and possibly look for factors that drive economic fluctuations.

4 Results

Table 1 shows estimates of Hamilton's model for quarterly real GNP growth over the 1952–1984 period. MSM0 is a replication of Hamilton's model, with parameter estimates that are close to, but not exactly the same as, those he reported. The difference is presumably due to a stricter convergence criteria for the maximum likelihood procedure that I used.[10]

The other models in Table 1 use data that has been revised to incorporate more recent information on national output, and implements a change in the official base year that is used to compute real GNP.[11] As a first test of robustness, the MSM1 set of parameters closely match MSM0 and capture conventional views about recessions and expansions. The α coefficients show that GNP tends to decline at a rate of 0.3 percent each quarter in regime 1, while growth each quarter averages a robust 1.1 percent in regime 2. It is also interesting

[8] Additional conditions for unboundedness are $\sigma_1 > 0$ and $q_{12} > 0$, with the latter being a necessary condition that obviously holds since period t' does occur. Continuity in the likelihood function means that interior solutions to the first-order conditions can never be global maxima.

[9] The test in Hansen (1992) only corrects the nuisance parameter problem and many researchers have ignored the fact that the singularity problem still exists.

[10] Following the data description in Hamilton's paper, real GNP growth was calculated using log-differences of data that start in 1951. Because the first four observations are saved for the initial autoregressive terms, the actual sample period starts in the second quarter of 1952 in all cases. Using a data and program printout that Hamilton provided, I verified that the difference in our results was not because of a programming error, and MSM0 was estimated using the same vintage of real GNP (which I pulled from an independent source). Also, Hamilton gave me a set of GAUSS programs for MSM estimation that he uses and distributes. They converged to parameters that were identical (to at least four decimal places) to those reported in Table 1. Hamilton originally estimated the model using a set of FORTRAN procedures.

[11] In 1989, the Department of Commerce changed the base year from 1982 to 1987. Recently, the National Income and Product Accounts were reformulated to emphasize an alternative measure for total real output, chain-weighted GDP, which solves many of the problems of changing base years. The business cycle characteristics of this series will be studied in a follow-up paper.

Table 1
Hamilton's MSM Specification for Real GNP Growth

	Data Sample					
	Hamilton	1952:2–1984:4 Revised			1952:2–1994:4	
	MSM0	MSM1	MSM2	MSM3	MSM4	MSM5
α_1	−0.3590	−0.3140	−1.5656	−2.0893	−1.5483	−2.1105
α_2	1.1635	1.1391	0.8475	0.8103	0.8057	0.7694
σ	0.7690	0.7648	0.8090	0.8332	0.7290	0.7589
ρ_1	0.0135	0.0427	0.3094	0.2738	0.3393	0.2887
ρ_2	−0.0575	−0.0594	0.1321	0.1989	0.1385	0.2142
ρ_3	−0.2470	−0.2341	−0.1441	−0.0997	−0.1552	−0.1007
ρ_4	−0.2129	−0.1205	−0.0904	−0.1520	−0.0523	−0.1302
q_{11}	0.7547	0.7520	0.2102	0.0000	0.2374	0.0000
q_{22}	0.9041	0.9086	0.9621	0.9741	0.9684	0.9801
Ln-lik	−181.26	−177.73	−175.61	−175.56	−210.27	−210.470

Notes: The underlined digits for MSM0 are the same as those in Hamilton (1989). (Hamilton reported: −0.3577, 1.1643, 0.7690, 0.0140, −0.059, −0.247, −0.213, 0.7550, 0.9049). For MSM3 and MSM5, the value of q_{11} converged to the boundary of 0, while the first-order conditions for the other parameters were satisfied.

that the serial correlation effects (ρ) are strongest at lags 3 and 4, and are decidedly negative. Conventional, linear AR(4) models for GNP growth show significant positive serial correlation at lag 1 and sometimes even lag 2. The main reason for the difference is that the transition probabilities of MSM1 capture most of the short-term persistence in the data. Also, the transition probabilities, $q_{11} = .75$ and $q_{22} = .91$, suggest important asymmetries in the business cycle where expansions tend to last considerably longer than recessions.[12]

However, with the revised data, two other local maximums, MSM2 and MSM3, were found in the 1952–1984 sample, and both have higher likelihood values. Table 1 shows that the quarterly growth rate of GNP in regime 1 tends to be below −1.5 percent in both cases, while growth in regime 2 is in the 0.8 percent range, which is only slightly above the entire sample's average (a single-regime concept). Also, the most significant ρ terms are for lags 1 and 2. Finally, both sets of parameters have $q_{11} < .50$. In fact, the model with the highest likelihood value, MSM3, has q_{11} at the 0 value boundary, which suggests that recessions last no longer than one quarter.[13]

A greater problem is that expanding the sample period to only 1986 removes the local peak with the plausible recession and expansion parameters that Hamilton reported. Table 1 reports results for the 1952–1994 period, and shows that the other parameter sets (MSM2 and MSM3) were robust to a change in the sample period. MSM4 closely matches MSM2, while MSM5 closely matches MSM3. Also, note that MSM4 with $q_{11} > 0$ has the highest likelihood value for the 1952–1994 period. The parameter set with $q_{11} = 0$ had the highest likelihood value in the shorter sample period.

Figure 1 provides additional information about these models. Here, I compare regime 1's full-sample probabilities for MSM1 to those from MSM4, using the entire 1952–1994 period.[14] The results show that MSM1 closely identifies regime 1 with NBER recession periods, while regime 1 for MSM4 does not fully capture the conventional recession/expansion dichotomy. Most interesting, only MSM1 matches the relatively mild 1990–1991 recession with a negative growth regime. In fact, the regime 1 probabilities from MSM5 (not

[12]For instance, the computed half-life of a recession is less than four quarters ($.7520^{2.4} = .50$), while the half-life of an expansion is over eight quarters ($.908607^{2.2} = .50$).

[13]The optimization routine converged to the $q_{11} = 0$ from various starting values, and this parameter set could be considered a corner solution. However, the first-order conditions for the other parameters were satisfied when q_{11} was held fixed at zero.

[14]Even though MSM1 was estimated for 1952–1984, the regime probabilities can still be calculated using a longer sample period. The results were almost completely immune to this change, and the MSM2 and MSM4 probabilities were also indistinguishable.

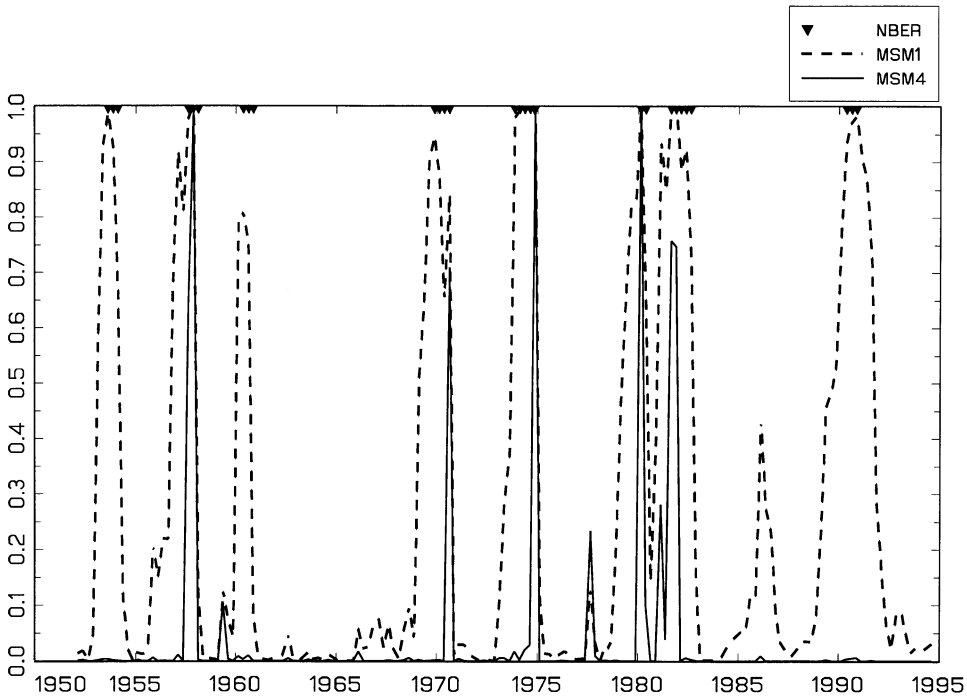

Figure 1
Full-sample probabilities for regime 1 from MSM1 and MSM4. Model parameters are in Table 1.

shown) are only above 50 percent in the three quarters that GNP fell more than 2 percent: 1958:1, 1974:1, and 1980:2.[15]

In Table 2, I show results from likelihood-ratio (LR) statistics for the cross-regime restrictions on σ and ρ that Hamilton used. In other words, I allow the equation variance and/or the serial correlation effects to differ across regimes.[16] To limit local maxima problems, the starting values for estimating these alternatives were based on the restricted MSM parameters reported in Table 1. This strategy produces a test statistic that has properties that resemble Lagrange multiplier tests in the sense that it is more conservative than allowing for a broad search across the set of permissible parameter values.

The LR test results for MSM1 reject the restriction on σ. In contrast, at least for this case, allowing ρ to vary across regimes does not lead to a significant improvement in the likelihood function. For MSM2, MSM3, and MSM5, the restriction on σ is less significant, while MSM4 shows hardly any improvement in fit. More significant are the differences in the importance of the ρ restrictions when the MSM2, MSM3, MSM4, and MSM5 parameters are used as starting values for the more general models. In some cases, allowing both σ and ρ to vary across regimes caused the ML procedure to get stuck at a very high likelihood value. Overparameterization is the problem. When q_{11} is low, separate ρ effects in regime 1 allow the model to fit almost perfectly the few observations where GNP declined over 1.5 percent. Specifically, when starting values such as MSM3 or MSM5 were used, the ML procedure led to parameter sets with both σ_1 and q_{11} close to zero and implausible ρ_1 coefficients.

In summary, simple MSMs with two regimes and four AR terms for capturing GNP growth dynamics are not robust across different sample periods, and show overparameterization problems. Next I explored

[15]In defining the most-likely regime dates, most researchers use a 50 percent cutoff value with the full-sample probabilities for each individual observation. Boldin (1992) discusses an alternative method that is an optimal joint-decision process. For these models, the differences are negligible.

[16]Here, I allow the ρ terms to depend on the current regime. Boldin (1992) and Hansen (1992) discuss and implement other ways to control for serial correlation. The most general model would allow ρ to vary according to the entire regime sequence (s_t, s_{t-1}, s_{t-2}, s_{t-3}, s_{t-4}). This type of model proves to be greatly overparameterized, however, and the alternative that I use is sufficient to point out problems with Hamilton's specification.

Table 2
Likelihood Ratio Tests for MSM1-5

	Ln-lik	LR	Prob
Sample Period: 1952–1994			
H_0. MSM1	−177.73		
H_A. $\sigma_1 \neq \sigma_2$	−175.06	5.33	0.021
H_B. $\rho_1 \neq \rho_2$	−176.61	2.23	0.692
H_C. $\sigma_1 \neq \sigma_2$, $\rho_1 \neq \rho_2$	−173.11	9.24	0.100
H_C vs. H_A	−173.53	3.07	0.546
H_C vs. H_B	−173.53	6.17	0.013
H_o. MSM2	−175.61		
H_A. $\sigma_1 \neq \sigma_2$	−174.16	2.90	0.089
H_B. $\rho_1 \neq \rho_2$	−173.19	4.84	0.304
H_C. $\sigma_1 \neq \sigma_2$, $\rho_1 \neq \rho_2$	−172.20	6.82	0.234
H_C vs. H_A	> −134.0	> 80.0	0.000
H_C vs. H_B	−173.02	0.33	0.562
H_o. MSM3	−175.56		
H_A. $\sigma_1 \neq \sigma_2$	−174.16	2.79	0.095
H_B. $\rho_1 \neq \rho_2$	−171.52	8.07	0.089
H_C. $\sigma_1 \neq \sigma_2$, $\rho_1 \neq \rho_2$	> −132.0	> 80.0	0.000
H_C vs. H_A	> −148.0	> 50.0	0.000
H_C vs. H_B	> −103.0	> 130.0	0.000
Sample Period: 1952–1994			
H_o. MSM4	−210.27		
H_A. $\sigma_1 \neq \sigma_2$	−210.24	0.06	0.810
H_B. $\rho_1 \neq \rho_2$	−207.97	4.61	0.330
H_C. $\sigma_1 \neq \sigma_2$, $\rho_1 \neq \rho_2$	−207.05	6.44	0.266
H_C vs. H_A	−207.05	6.39	0.172
H_C vs. H_B	−207.05	1.83	0.176
H_o. MSM5	−210.47		
H_A. $\sigma_1 \neq \sigma_2$	−209.45	2.03	0.153
H_B. $\rho_1 \neq \rho_2$	−205.16	10.61	0.031
H_C. $\sigma_1 \neq \sigma_2$, $\rho_1 \neq \rho_2$	> −167.0	> 86.0	0.000
H_C vs. H_A	> −167.0	> 85.0	0.000
H_C vs. H_B	> −140.0	> 160.0	0.000

Notes: LR statistics are two times the difference between the alternative and null model log-likelihood values. For the first three tests in each panel, the null model is H_0. In all cases, the null (simpler) model's parameters were used as the starting values in the ML routine that computed the reported log-likelihood value for the alternative model. In the ">" cases, the ML routine did not converge after 200 iterations, and the 200th iteration's log-likelihood value exceeded the number reported.

modifications that relaxed the cross-regime restrictions, reduced the AR order, and made the switching process more plausible. First, relaxing the cross-regime restrictions on σ and ρ, even with a lower number of AR terms in each equation, was almost always a (seemingly) statistically significant improvement even when σ_1 and q_{11} were not close to zero. Second, assuming that all regimes last more than one quarter invariably increased the likelihood function without requiring more parameters to be estimated.

Table 3 reports a model, MSM6, with no AR terms and no restrictions on σ, that was estimated for the 1952–1994 period. In this case, a third-order transition process with the following properties was used:

1. Each regime must last at least two quarters; i.e., prob($s_t = i \mid s_{t-1} = i, s_{t-2} \neq i$) = 1 and of course prob($s_t \neq i \mid s_{t-1} = i, s_{t-2} \neq i$) = 0.

Table 3
Alternative Two-Regime Switching Models

	MSM6		MSM7		MSM8	
Regime:	1	2	1	2	1	2
α	−0.3474	0.9882	−1.5749	0.7954	0.4432	0.9818
se	0.2328	0.0783	0.2233	0.0962	0.1433	0.0746
σ	0.9139	0.7031	0.5961	0.7466	0.8041	0.6221
se	0.0827	0.0487	0.1423	0.0419	0.0645	0.0685
$\rho 1$			−0.7219	0.2780	0.3268	0.1761
se			0.2281	0.0712	0.0835	0.0883
$\rho 2$			0.0780	0.1105	0.5527	−0.3502
se			0.3344	0.0720	0.0964	0.0813
q_{211}	1.0		0.0		0.5350	
					0.1223	
q_{111}	0.5850		0.0		0.7238	
se	0.1884				0.0791	
q_{221}		0.9552		1.0		1.0
se		0.1571				
q_{222}		0.9348		0.9750		0.8985
se		0.0273		0.0124		0.2273
Ln-Lik	−215.0539		−208.4738		−206.946	
LR test of $\sigma_1 \neq \sigma_2$			13.16		16.20	
prob			0.011		0.003	
LR test of $\rho_1 \neq \rho_2$			9.22		12.26	
prob			0.027		0.007	

Notes: 1952:2–1994 is the sample period for all models. The presentation format is different than Table 1. The coefficients for each regime are in separate columns. The numbers in the "se" rows are ML parameter standard errors calculated by numerically approximating the log-likelihood function's Hessian. When the q parameters converged to a boundary value (0.0 or 1.0), this condition was imposed for computing the standard errors for the other parameters.

2. Prob($s_t = i \mid s_{t-1} = i, s_{t-2} = i, s_{t-3} = i$) ≠ prob($s_t = i \mid s_{t-1} = i, s_{t-2} = i, s_{t-3} \neq i$) is allowed.[17]

On the plus side, the estimated parameters for MSM6 have plausible business-cycle interpretations, they were robust across different sample periods and variations in the assumptions about the transition process, and a second local maximum was not found. However, the log-likelihood fit is considerably worse than the restricted AR(4) models in Table 1 and the unrestricted cases reported in Table 2. The two other models in Table 3, MSM7 and MSM8, are AR(2) specifications that correspond to distinct local maximums and have log-likelihood fits that reject both the nested AR(0) specification, MSM6, and models that restrict the σ and ρ parameters in the same manner as Hamilton's model.[18] While the parameters in MSM7 are similar in many respects to those in MSM4, the MSM8 coefficients are not easily reconciled with conventional business-cycle beliefs. Figure 2 graphs the regime 1 probabilities for MSM6 and MSM7, and we see some correspondence to conventional business-cycle dates. However, Figure 3 shows that MSM8 does not capture business-cycle patterns at all, which is disconcerting since it has a higher likelihood value than both MSM4 and MSM7.

[17]Technically, the switching process is still Markovian since it is not time-varying and this model can set up as a first-order process with six regimes that are defined by the admissible sequences of the original regime labels and cross-regime restrictions on α and σ.

[18]Because of the numerous local-maxima problem, the parameter standard errors in Table 3 are probably not reliable for computing confidence intervals. Instead, they are reported because they provide some useful information about the (local) curvature of the likelihood function.

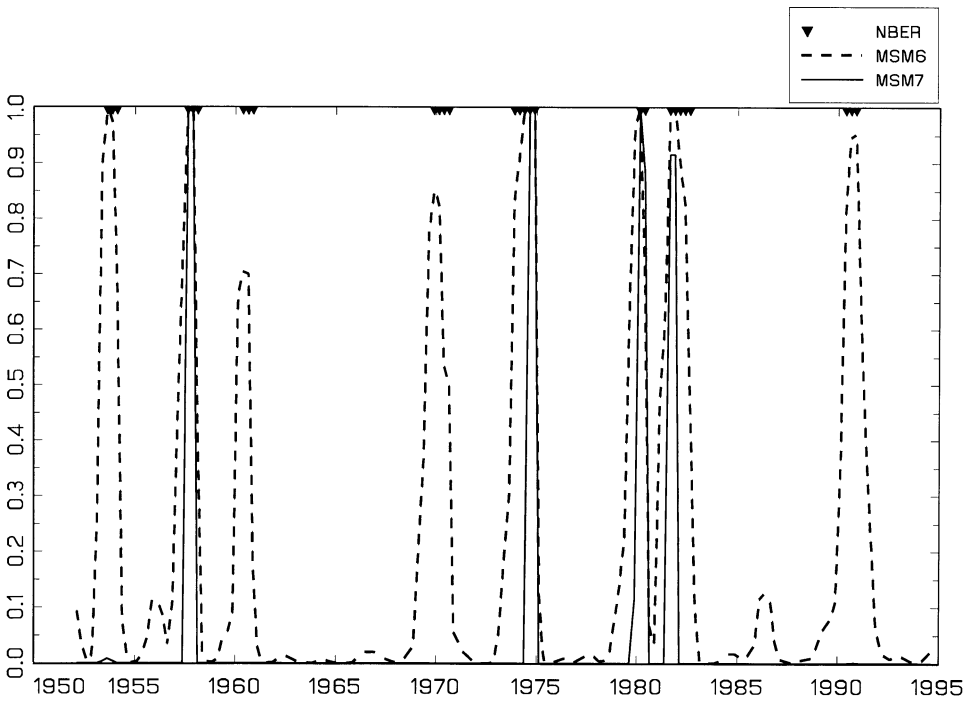

Figure 2
Full-sample probabilities for regime 1 from MSM6 and MSM7. Model parameters are in Table 3.

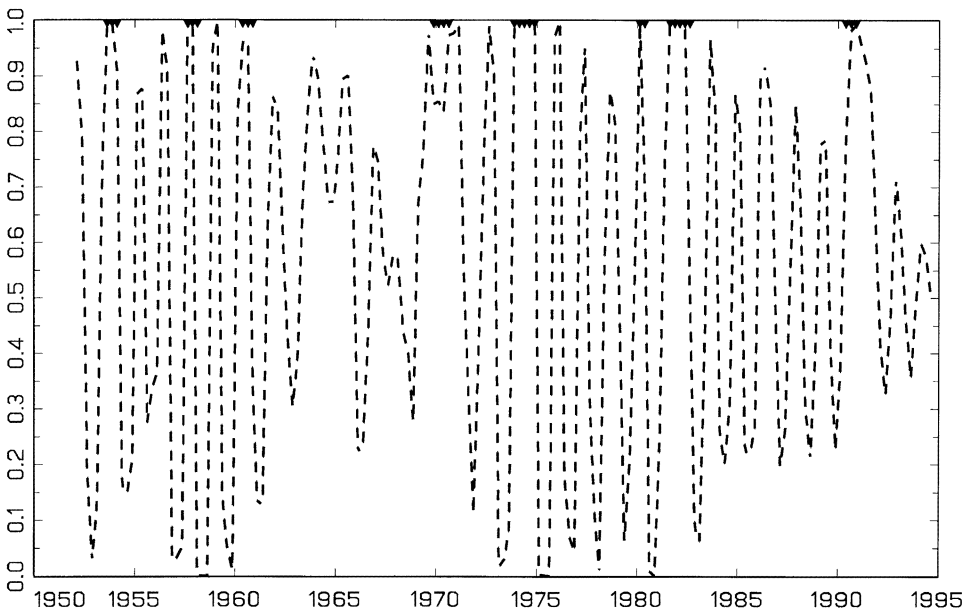

Figure 3
Full-sample probabilities for regime 1 from MSM8. Model parameters are in Table 3.

Table 4
Three-Regime Switching Model MSM9

Regime:	1	2	3
α	−0.4792	1.7016	0.6598
se	0.2750	0.0849	0.0815
σ	0.9620	0.4740	0.5558
se	0.0873	0.0591	0.0490
$\rho 1$	0.1834	0.4297	0.2707
se	0.1196	0.0988	0.1024
$\rho 2$	0.6266	−0.2972	0.1043
se	0.1579	0.0889	0.0916
q_{211}	0.6384		
se	0.1705		
q_{111}	0.7117		
se	0.0966		
q_{122}		1.0	
se			
q_{222}		0.4984	
se		0.0726	
q_{233}			0.1024
se			0.0355
q_{333}			0.1024
			0.0355
Ln-Lik		−200.6995	
LR test of $\sigma_1 \neq \sigma_2$,		13.14	
prob		0.040	
LR test of $\rho_1 \neq \rho_2$		22.81	
prob		0.001	

Notes: 1952:2–1994 is the sample period. The restriction $q_{233} = q_{333}$ was used for this model. See Table 3 notes for further information.

The most interesting results came from a modified MSM with three regimes and two AR terms. In Table 4, MSM9 has a plausible negative growth rate for regime 1 that can be identified with recessions. Regime 2 covers post-recession rapid-recover periods, and regime 3 shows more moderate growth for the remainder of the expansion. Important and plausible variation across regimes for the σ and ρ coefficients are also seen. Formal likelihood ratio tests in the table suggest that the σ and ρ parameters should not be held constant across regimes, and at least two AR terms are significant.

It must be noted that the regime-switching process for MSM9 is even more complicated than the models in Table 3 and incorporates some important restrictions. Not only must each regime be longer than one quarter, regime 1 can only lead to regime 2, which only leads to regime 3, which only leads back to regime 1. In other words, the regime progression is unidirectional. Also, the ML procedure converged to a case where regime 2 always lasts at least three quarters. Another interesting result is that the odds of remaining in regime 1 actually increases after two quarters. Finally, to get sensible parameters, I had to make a strong restriction on the transition process for regime 3. In this specification, the probability of remaining in regime 3 is constant after two quarters. If this restriction is not made, the ML procedure tends to converge to a case where $\alpha_1 > 0$, the σ value for regime 3 is implausibly low, and the odds of leaving regime 3 are over 90 percent after two

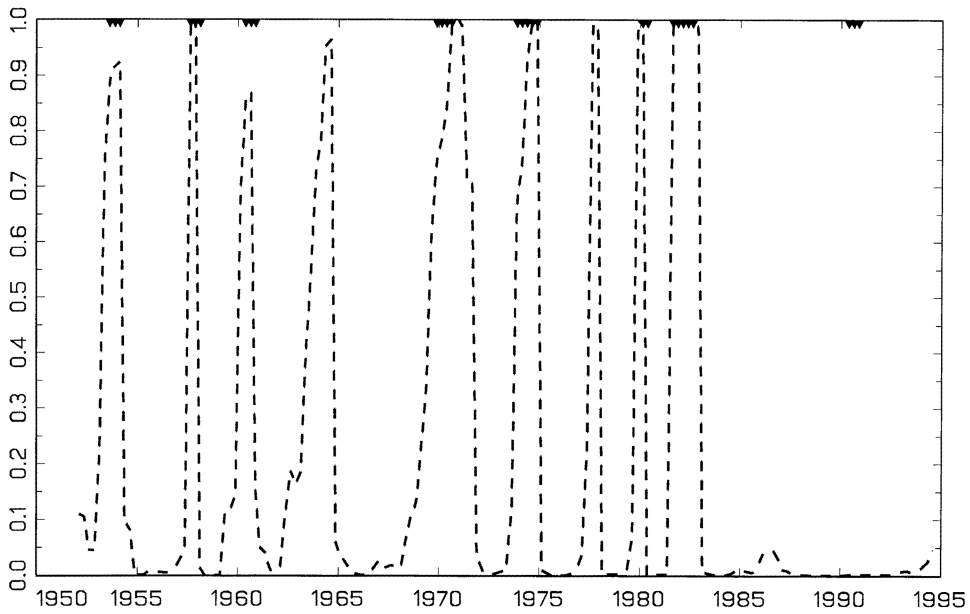

Figure 4
Full-sample probabilities for regime 1 from MSM9. Model parameters are in Table 4.

periods.[19] In essence, the transition restrictions that I used require the total duration of an expansion (combining regime 2 and 3) to be over a year. This restriction yields an estimate for the odds of a recession starting at around 10 percent each quarter.

Figure 4 shows the corresponding full-sample probabilities for regime 1; it is not surprising, given the parameter values, that the NBER recession dates are captured fairly well, although other local maxima were found for this exact specification and one case consistently had a higher likelihood value. However, the improvement in the likelihood value relies on implausibly low σ values for some regimes. More important, the general characteristics of MSM9 in Table 4 were robust across changes in the sample periods. In other words, a local maxima near the parameter set that I reported was not lost when different estimation periods were used. For instance, with 1952–1984 as the sample period, the three α coefficients were −0.2756, 1.6617, and 0.8477; with 1960–1994 as the sample period, they were −0.5147, 1.8269, and 0.5896; and with 1960–1989 as the sample period, they were −0.5864, 1.8075, and 0.5784. These values are reasonably close to the three distinct mean growth rates for MSM9 of −0.4792, 1.7016, and 0.6598 percent each quarter. The other parameters were similar as well.

5 Future Research

Similar multiple local maxima and overparameterization problems can be shown with various output and employment series, and seem to be inherent in efforts to capture business-cycle patterns with MSMs using ML estimation. It may then be tempting to rely on other estimation procedures that have been proposed, namely, EM algorithms (Hamilton 1990) and Gibbs sampling (Albert and Chib 1993) techniques. But these cannot be true solutions since they are also based on "maximizing the likelihood" concepts as well. One possible solution is to follow Hamilton (1991), which proposes a Bayesian adjustment to the likelihood function that accounts for prior beliefs about the parameters. However, there will always be a concern that the results depend on the priors and not the data being analyzed. The better course seems to be to search for flexible, yet parsimonious specifications such as MSM9 that limit excessive regime switching.

[19]The log-likelihood value for the alternative model is over −197.0, and would be considered a statistically significant improvement according to conventional LR test criteria, however.

Given the plausibility of the three-regime switching model characteristics, a review of the consistent findings from linear time-series analysis that output is not trend reverting is warranted. Lam (1990) has already presented evidence of less significant output-shock persistence in a two-regime framework. It is likely that within a particular regime, especially the normal expansion phase (regime 3), trend reversion is significantly stronger than what is implied by a single-regime model. Reliable tests for the number of regimes will be needed before results from multiple-regime models are taken seriously, however.

6 Conclusion

This article highlighted the extent to which multiple local maxima problems plague seemingly reasonable MSM specifications for output dynamics. Making a positive contribution, I demonstrated that a three-regime switching model is much more useful for the analysis of business cycles than the specification that Hamilton first proposed. The results show even greater asymmetries and nonlinear characteristics in GNP growth. Nonetheless, the general insight and value of Hamilton's multiple-regime framework is retained.

References

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

Boldin, M. (1990), "Characterizing Business Cycles with a Markov Switching Model: Evidence of Multiple Equilibria," Federal Reserve Bank of New York Research Paper 9037.

Boldin, M. (1992), "Using Switching Models to Study Business Cycle Asymmetries: Overview of Methodology and Applications," Federal Reserve Bank of New York Research Paper 1992-11.

Cosslett, S. and L. Lee (1984), "Serial Correlation in Discrete Variable Models," *Journal of Econometrics* 85, 79–97.

Filardo, A. (1994), "Business-Cycle Phases and Their Transitional Dynamics," *Journal of Business and Economic Statistics* 9, 299–308.

Goodwin, T. (1993), "Business-Cycle Analysis with a Markov-Switching Model," *Journal of Business and Economic Statistics* 9, 331–339.

Hamilton, J. (1989), "A New Approach to Economic Analysis of Nonstationary Time Series," *Econometrica* 57, 357–384.

Hamilton, J. (1990), "Analysis of Time Series Subject to Changes in Regime," *Journal of Econometrics* 45, 39–70.

Hamilton, J. (1991), "A Quasi-Bayesian Approach to Estimating Parameters for Mixtures of Normal Distributions," *Journal of Business and Economic Statistics* 9, 27–40.

Hansen, B. (1992), "The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model," *Journal of Applied Econometrics* 7, S61–S82.

Lam, P.-K. (1990), "The Hamilton Model with a General Autoregressive Component," *Journal of Monetary Economics* 26, 409–432.

Advisory Panel

Jess Benhabib, New York University

William A. Brock, University of Wisconsin-Madison

Jean-Michel Grandmont, CEPREMAP-France

Jose Scheinkman, University of Chicago

Halbert White, University of California-San Diego

Editorial Board

Bruce Mizrach (editor), Rutgers University

Michele Boldrin, University of Carlos III

Tim Bollerslev, University of Virginia

Carl Chiarella, University of Technology-Sydney

W. Davis Dechert, University of Houston

Paul De Grauwe, KU Leuven

David A. Hsieh, Duke University

Kenneth F. Kroner, BZW Barclays Global Investors

Blake LeBaron, University of Wisconsin-Madison

Stefan Mittnik, University of Kiel

Luigi Montrucchio, University of Turin

Kazuo Nishimura, Kyoto University

James Ramsey, New York University

Pietro Reichlin, Rome University

Timo Terasvirta, Stockholm School of Economics

Ruey Tsay, University of Chicago

Stanley E. Zin, Carnegie-Mellon University

Editorial Policy

The *SNDE* is formed in recognition that advances in statistics and dynamical systems theory may increase our understanding of economic and financial markets. The journal will seek both theoretical and applied papers that characterize and motivate nonlinear phenomena. Researchers will be encouraged to assist replication of empirical results by providing copies of data and programs online. Algorithms and rapid communications will also be published.

ISSN 1081-1826