

Erlandsson, Ulf

Working Paper

Reconnecting the Markov Switching Model with Economic Fundamentals

Working Paper, No. 2004:4

Provided in Cooperation with:

Department of Economics, School of Economics and Management, Lund University

Suggested Citation: Erlandsson, Ulf (2004) : Reconnecting the Markov Switching Model with Economic Fundamentals, Working Paper, No. 2004:4, Lund University, School of Economics and Management, Department of Economics, Lund

This Version is available at:

<https://hdl.handle.net/10419/259880>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Reconnecting the Markov switching model with economic fundamentals

Revised version

ULF G. ERLANDSSON*

November 4, 2004

Abstract

This paper seeks to investigate and remedy the apparent inability of Markov regime switching models to predict future states in the medium to long term. We show that projected time varying transition probability series in the model may be biased towards predicting regime switches with high probability in the short run, and as a consequence it is hard or impossible to obtain longer run inference. We propose a penalized maximum likelihood estimator where non-smoothness in the transition series has negative influence on the likelihood function, which is shown to remedy the short run bias. In an empirical investigation of U.S. real GDP, the penalized model works better in terms of forecasting future recessions as defined by the NBER business cycle dating.

KEYWORDS: regime switching, transition probability, forecasting

JEL CLASSIFICATION: C13; C32; E32

*Department of Economics, Lund University; P.O. Box 7082; 220 07 LUND; email: ulf.erlandsson@nek.lu.se; home-page: www.nek.lu.se/nekuer

Introduction

Few non-linear time series models have attained the same level of popularity as the Markov regime switching model of Hamilton (1989). The division of economic variables into different states, such as contraction and expansions phases of GDP, enjoys a intuitive advantage over more continuous models. A very large part of contemporary research focuses on improving the Hamilton model by including various extensions from other parts of the time-series literature in order to match the moments in the data.

A less probed but equally important area emphasizes the other novel invention following the Hamilton model: the determination of states. We may very well match moments of the model perfectly but if we assume an incorrect or non-optimal structure to determine the probability of states, the performance of the model will suffer. In the plain vanilla Hamilton model, the transition between states is governed by a constant probability. This assumption is rarely questioned, although intuitively we accept that the probability that a bull-market continues is higher after only a few months of rising stock prices than when the bull-market has persisted for a longer period of time. To reflect this in the model, one may let the transition probability be time varying. The first published paper to recognize this possibility is Diebold, Lee and Weinbach (1994).

A number of papers have applied the Diebold et al. methodology: Gray (1996) for interest rates, Tronzanon, Psaradakis and Sola (2003) for target zones and Abiad (2003) in the currency crisis context. Results have been mixed, and it seems that one of the main reasons for the relatively infrequent use of this approach is the difficulty to obtain sensible parameter estimates. The problems are exacerbated in the multivariate transition equation setting. Statistical inference in this environment can be very difficult, since the model in essence tries to estimate a model with two unobserved variables. As always in the Markov regime switching (henceforth denoted MS) model, we try to estimate the unobserved state variable. Moreover, we seek the relation between exogenous variables and the dependent variable in the transition equation which also is unobserved.

This paper seeks to establish that the maximum likelihood estimator will be biased toward finding a parameterization of the model that leads to a projected transition probability series with very abrupt shifts. In general, a parameterization that produces a high probability to switch regimes a short time interval prior to an observed shift will be preferred to a parameterization with a lower probability to switch but during a longer time period, irrespective of the true data generating process.

Although the bias reflects an optimal moment-matching, the use of the model becomes deeply restricted due to this bias. For policy purposes, parameter values may be of lesser importance than the ability to obtain early warnings of an upcoming switch. The exact size of the contraction/expansions trends in a model of GDP are relatively unimportant to a policy maker who intends to take measures that only have lagged effects. The average time before an interest rate change gives

effect is estimated to be around 12-24 months. For a central banker, early warnings of the type "with a 50% probability, GDP growth will enter the contraction phase within the next 18 months" will be preferred to a statement of the type "with 95% probability, we will enter the contraction phase next month."

We provide a possible solution to the short run bias problem by introducing a penalty term in the log likelihood function. It penalizes non-smooth behavior of the transition probabilities series with a weight chosen by the researcher. In that sense, it is a subjective approach, but our results indicate that this addition remedies a number of problems inherent in the standard estimation procedure. Foremost, it increases the correlation between the projected transition probability series and the true data-generating transition probability process.

In an empirical exercise, we investigate U.S. real GDP in a regime switching setting. The penalized maximum likelihood method finds a different and much smaller set of variables predicting switches to recession to include in the final model than the ordinary maximum likelihood does. When forecasting the NBER recessions, the penalized model exhibits better performance than the benchmark models for horizons exceeding 4 quarters, both in- and out-of-sample.

In section 2, we introduce the baseline model and the maximum likelihood estimation procedure. Thereafter we discuss MS models with time varying transition probabilities (TVP) as a bounded model similar to the the popular Threshold Autoregression family models. Section 4 is dedicated to a proof in a simple model that the maximum likelihood estimator will select a short-run variable prior to a longer run one, irrespectively of the underlying data-generating process (DGP). In section 5, simulation evidence corroborates these results in the stochastic setting. We propose a remedy to the short run bias as well, and analyze the effects of using a penalty term in various settings using Monte Carlo analysis in section 6. We apply the proposed method on actual data in section 7. Section 8 concludes.

Model and Estimation Procedure

We base the discussion on a simple form of the Hamilton (1990) Markov regime switching model. The baseline model with constant transition probabilities:

$$\Delta y_t = \mu_{S_t} + \epsilon_t \quad (1)$$

where S_t is a state variable that follows a first order Markov chain with transition probability matrix:

$$\mathbf{P} = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix} \quad (2)$$

where, in turn, p_{ij} denotes the probability to go from state i to state j . The iterative procedure to estimate this kind of model is presented in Hamilton (1994).

The number of extensions made to this simple model, and the combinations thereof, can be counted in the hundreds. Most of these seek to engineer to model in

a way as to have a better fit to the data, modifying the elements of equation 1, e.g. by introducing exogenous variables, auto-regressive parameters and ARCH effects. A smaller number of studies, e.g. Diebold, Lee and Weinbach (1994) , have focused on modelling the probability to switch to other regimes, as in equation 2, noting that \mathbf{P} by no means have to be constant. In the general 2 state case:

$$\mathbf{P}_t(\mathbf{Z}_t) = \begin{bmatrix} g(\mathbf{Z}_t^1) & 1 - g(\mathbf{Z}_t^1) \\ 1 - h(\mathbf{Z}_t^2) & h(\mathbf{Z}_t^2) \end{bmatrix} \quad (3)$$

where $g, h \rightarrow [0, 1]$.¹ This will be referred to the time-varying transition probability (TVP) model. The functional form of f, g is usually chosen to be of probit or logit type. We will assume the logit style functional form for both h, g such that:

$$h(x) = g(x) = \frac{\exp(x)}{1 + \exp(x)} = f(x) \quad (4)$$

In order to estimate the Hamilton Markov regime switching model we iterate on the equations:

$$\xi_{t|t} = \frac{\xi_{t|t-1} \odot \eta_t}{\mathbf{1}'(\xi_{t|t-1} \odot \eta_t)} \quad (5)$$

and

$$\xi_{t|t-1} = \mathbf{P}_{t-1}(\mathbf{Z}_{t-1})' \cdot \xi_{t-1|t-1} \quad (6)$$

where η_t is a $(N \times T)$ matrix of each N states conditional density based on the parameter vector θ . For the 2 state case:

$$\eta_t = \begin{bmatrix} \frac{1}{\sqrt{2\pi}\sigma_1} \exp\left\{-\frac{(y_t - \mu_1)^2}{2\sigma_1^2}\right\} \\ \frac{1}{\sqrt{2\pi}\sigma_2} \exp\left\{-\frac{(y_t - \mu_2)^2}{2\sigma_2^2}\right\} \end{bmatrix} \quad (7)$$

The log likelihood to be maximized is given by:

$$L(\theta) = \sum_{t=1}^T \log \mathbf{1}'(\xi_{t|t-1} \odot \eta_t) \quad (8)$$

A number of other estimation methods are available, see e.g. Filardo and Gordon (1998) for a Gibbs sampling approach in the time varying transition probability context.

¹Extending the number of states to N is straightforward in theory, but harder in practice since the number of coefficients grows with a factor of $2N$ rendering estimation difficult.

Bounded Regime Switching Processes

To be able to conduct simulation exercises, we now introduce a regime switching parameterization where the projected time-series is only dependent upon the parameter vector and a single vector of random disturbances.

The process will be based on the previously considered model in equation 1:

$$\Delta y_t = \mu_{S_t} + \epsilon_t \Rightarrow y_t = y_{t-1} + \mu_{S_t} + \epsilon_t \quad (9)$$

The long run drift of y_t can be calculated using the ergodic (unconditional) probabilities:

$$P(S_t = j) = \pi = \left[\begin{array}{c} (1 - p_{22})/(2 - p_{11} - p_{22}) \\ (1 - p_{11})/(2 - p_{11} - p_{22}) \end{array} \right] \quad (10)$$

which can be obtained by solving the eigenvalue problem $|\mathbf{P} - \lambda \mathbf{I}_2| = 0$. Explicitly, the long drift in the N state model becomes

$$\bar{\mu} = \sum_{j=1}^N P(S_t = j) \cdot \mu_j$$

Using the long term drift, it easy to see that in the long-run, this process will mimic a random walk with drift. For many purposes, however, it seems unreasonable that a variable - in the long run - should follow such a process. Examples could be trade balance, debt-to-GDP ratios and real exchange rates. One could also use the same argument to build a model of financial bubbles as in Schaller and van Norden (2002). It is likely to be some reversion back to some, yet undefined, mean once we reach a level that is much higher or lower than the posited mean. Assume that $\mu_2 = -\mu_1$ and $\mu_1 \geq 0$. Moreover, assume that there is a bound $a \geq 0$ so that $P(S_{t+1} = 2 | y_t \geq a) = 1$ and $P(S_{t+1} = 1 | y_t \leq -a) = 1$. In words, if the level process y_t exceeds/goes below $a/-a$, we automatically switch back to a state that reverts the process in the other direction. This is analogously to a Threshold Auto-Regressive (TAR) model. Another way to express this is that the boundary model in effect has time varying transition probabilities. The two transition matrices are:

$$\mathbf{P}_t = \left\{ \begin{array}{ll} \left[\begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array} \right] & \text{if } \|y_t\| > a \\ \left[\begin{array}{cc} p_{11} & p_{12} \\ p_{21} & p_{22} \end{array} \right] & \text{if } \|y_t\| \leq a \end{array} \right. \quad (11)$$

The long run mean \bar{y} of the bounded process can be calculated but most noteworthy here is the fact that the process y_t will never be 'far' away from its long-run mean, which makes an argument for the variable to posses a form of stationarity. The conditions for stationarity of this process can be found in Karlsen (1990).

An even more flexible and sensible,² version of the bounded model is where the

²The discrete boundaries previously considered cannot be estimated with gradient based optimization algorithms. As in the TAR case, one usually would resort to grid-based estimation procedures.

probability to go to the reversion state is dependent upon the distance the process is from its long run mean. Hence, we would posit, with $\|\cdot\|$ denoting an appropriate metric, that $P(S_{t+1} = j | S_t = i; \|y_t - \bar{y}_t\|)$ is large when $\|y_t - \bar{y}_t\|$ is. We can translate this to the more general form of the transition matrix in equation 3:

$$\mathbf{P}_t(\beta; \|y_t - \bar{y}_t\|) = \begin{bmatrix} f(\beta_1 \cdot \|y_t - \bar{y}_t\|) & 1 - f(\beta_1 \cdot \|y_t - \bar{y}_t\|) \\ 1 - f(\beta_2 \cdot \|y_t - \bar{y}_t\|) & f(\beta_2 \cdot \|y_t - \bar{y}_t\|) \end{bmatrix} \quad (12)$$

By setting $\beta_1 \neq \beta_2$, the boundaries are allowed to be asymmetric. Once the process y_t goes 'far' off from the long run mean, the probability that we will switch to the state where we revert back increases. In the limit, this switch will happen with probability 1. Hence, the same argument for global stationarity as in the fixed threshold setting applies. We also note that the process does not have to have both an upper and lower bound in order to be globally stationary. If the long run drift term is positive/negative, the process will be globally stationary if there is an upper/lower bound.

Using the baseline model (1) with the associated transition matrix in (12), we see that we have a process that is only a function of the parameter vector $\theta \in \{\mu, \beta\}$ and the vector of disturbances ϵ . Besides the advantage of being parsimonious, this model can function as a vehicle for simulation studies, since artificial data is easily generated.

Regime switches and short-run indicators

Now we will consider a simple case of the TVP model which brings about a paradox with serious economic implications. Consider the growth of the debt-to-GDP ratio. We assume that there are two states: the first where debt is growing, and a second where debt is decreasing. Economic theory suggest that there is an upper bound to how much debt in relation to income can grow, since rational lenders will not supply more credit once debt-to-GDP reaches an 'unsustainable' level. The point in time where credit dries up due to unsustainability will be depicted as a regime switch, where the ratio cannot rise anymore,³ but switches to the decreasing state as the government is forced by creditors to impose policies to this end. There will be some heterogeneity in opinion of what is a sustainable level; hence the exact timing of the switch is not predictable, although the probability to switch is. One motivation for the sudden shift is that there seems to be herding effects. Once a large investor stops buying a certain country's debt, other investors tend to follow to avoid a liquidity squeeze.

To estimate such a model on empirical data, we would use the set-up (state 1 corresponds to the rising debt state):

$$\Delta y_t = \mu_{S_t} + \epsilon_t \quad (13)$$

³Of course, GDP can decrease with the effect of the ratio rising, but this is not likely to happen for extended periods of time. We will not consider that case here.

with

$$\mathbf{P}_t(\theta; y_{t-1}) = \begin{bmatrix} f(\alpha_1 + \beta_1 \cdot y_{t-1}) & 1 - f(\alpha_1 + \beta_1 \cdot y_{t-1}) \\ p_{21} & p_{22} \end{bmatrix}$$

where Δy_t is the growth of the debt-to-GDP ratio. In this case, it seems unlikely that there is a long-run drift of the dependent variable, so $\bar{\mu}$ in equation (12) is set to zero.

In empirical work, however, we do not have the luxury of knowing exactly what variables to include in the transition equations, but have to discriminate between a number of possible candidates. Besides using the dependent variable itself, let us assume we also observe a binary variable called x_t that takes on the value 1 one unit of time prior to the crisis and is zero otherwise. In our setting, an example of this could be a negative change in the credit rating of the country's debt. The causality of this in relation of the probability to switch to the credit constrained regime is ambiguous. On one hand, the change is induced by the gradual rise in the probability to enter the credit constrained state. On the other hand, a worsening in the credit rating could be argued to lead to a rising probability to switch to the constrained regime, *ceteris paribus*. So one could argue that both variables - i.e. the total debt-to-GDP and the change in credit rating - should be included in the model. The former variable depicts long-run fundamentals and the latter short-run sentiments.

Hence, we would like to estimate the model with the following transition probability parameterization:

$$\mathbf{P}_t(\theta; y_{t-1}) = \begin{bmatrix} f(\alpha_1 + \beta_1 \cdot y_{t-1} + \beta_2 \cdot x_{t-1}) & 1 - f(\alpha_1 + \beta_1 \cdot y_{t-1} + \beta_2 \cdot x_{t-1}) \\ p_{21} & p_{22} \end{bmatrix}$$

In this model, we will find that β_2 is very significant and β_1 insignificant, even if the data is generated using the model the model in (13)! The reason for this will be shown below. At this stage, we want to note that for policy purposes, a model that gives us as much advance warning of an oncoming debt crisis as possible will be preferred to one that gives us very little time to react. But paradoxically the best econometric fit is obtained with a model that is more or less worthless for policy purposes since it only gives advance warning in the period prior to the crisis.

To see why we by traditional econometric criteria will select the variable with the short duration prior to the shift, we should observe the likelihood function. Suppose there occurs a regime shift at time T and that $S_t = 1$ for $t = 1, 2, \dots, T-2, T-1$. We have two binary possible candidates: x^A that produces a low probability $(1 - \pi_1) = 1 - f(x^A)$ to switch to regime 2 from time $T - j : T - 1$, $j > 1$; and x^B that produces a very high probability $(1 - \pi_2) = 1 - f(x^B)$ to switch at time $T - 1$ but a 0 probability otherwise, and $(1 - \pi_1) < (1 - \pi_2) \Rightarrow \pi_1 > \pi_2$. The corresponding transition matrices are

(A)

$$\mathbf{P}_{T-j:T-2}^A = \begin{bmatrix} \pi_1 & 1 - \pi_1 \\ p_{21} & p_{22} \end{bmatrix} \Rightarrow \mathbf{P}_{T-1}^A = \begin{bmatrix} \pi_1 & 1 - \pi_1 \\ p_{21} & p_{22} \end{bmatrix}$$

and
(B)

$$\mathbf{P}_{T-j:T-2}^B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{P}_{T-1}^B = \begin{bmatrix} \pi_2 & 1 - \pi_2 \\ p_{21} & p_{22} \end{bmatrix}$$

We furthermore assume that we are certain to have been in regime one for the whole period prior to the switch so that $\xi_{t|t} = [1 \ 0]'$ for $t = 1, 2, \dots, T-2, T-1$. Because of that, if we disregard the mean zero random part, we can simplify η_t in (7) to

$$\eta_t = [a \ b]'$$

for $t = 1 \dots T-1$ and we also have that $a > b > 0$. For the time $t = T \dots T+k$, we have

$$\eta_t = [b \ a]'$$

Just looking at the time $T-1$ transition matrices, (B) gives us a better explanation of the dependent variable than (A). To verify this, we note that at time T , the inner part of the likelihood functions is:

$$L_T = \mathbf{1}' \left(\xi_{T|T-1} \odot \eta_T \right) = \mathbf{1}' \left(\mathbf{P}_{T-1}' \xi_{T-1|T-1} \odot \eta_T \right) \quad (14)$$

For the case (A) and (B) this reduces to: $\underbrace{\pi_1 b + (1 - \pi_1)a}_{L_T^A}$ and $\underbrace{\pi_2 b + (1 - \pi_2)a}_{L_T^B}$.

Setting these equal and solving yields:

$$\underbrace{\pi_1}_{+} \underbrace{(b-a)}_{-} = \underbrace{\pi_2}_{+} \underbrace{(b-a)}_{-}$$

so that

$$L_T^A < L_T^B$$

It is straightforward to verify that also $L_{T-j:T-1}^A < L_{T-j:T-1}^B$. Hence, an ordinary maximum likelihood estimator would prefer variable (B) to variable (A) in this setting. So far, this is not controversial.

But what happens if we try to estimate a model where both x^A, x^B are included in the transition equations? To do this we need to elaborate on the relation between the transition matrix and the functions producing it. Consider that $\pi_1 = f(\alpha_1 - \beta_1 x^A)$, so that a higher value of β_1 means a higher probability to switch regimes, and $\pi_2 = f(\alpha_2 - \beta_2 x^B)$. We assume that both x variables are positive. We first note that a change of β_1 has effects on two likelihood elements. The first is at time $T-1$:

$$\frac{\partial L_{T-1}^A}{\partial \beta_1} = \frac{-x_{T-1}^A \exp(\alpha_1 - \beta_1 x_{T-1}^A)}{\left[1 + \exp(\alpha_1 - \beta_1 x_{T-1}^A)\right]^2} \cdot (a - b) < 0 \quad (15)$$

and

$$\frac{\partial L_{T-1}^B}{\partial \beta_2} = 0 \quad (16)$$

In essence, equation (15) means that if there is a probability for a regime switch but none occurs, this affects the likelihood value negatively irrespective of the true underlying transition probability process.

The derivative of L_{T-1}^B has been calculated using the fact that $x_{T-1}^B = 0$. Looking at time T derivatives instead we obtain:

$$\frac{\partial L_T^A}{\partial \beta_1} = \frac{-x_T^A \exp(\alpha_1 - \beta_1 x_T^A)}{[1 + \exp(\alpha_1 - \beta_1 x_T^A)]^2} \cdot (b - a) > 0 \quad (17)$$

and

$$\frac{\partial L_T^B}{\partial \beta_2} = \frac{-x_T^B \exp(\alpha_2 - \beta_2 x_T^B)}{[1 + \exp(\alpha_2 - \beta_2 x_T^B)]^2} \cdot (b - a) > 0 \quad (18)$$

Looking at the global likelihood, using equations (15)-(18), we obtain:

$$\frac{\partial L^A}{\partial \beta_1} = \frac{\partial L_{T-1}^A}{\partial \beta_1} + \frac{\partial L_T^A}{\partial \beta_1}$$

and

$$\frac{\partial L^B}{\partial \beta_2} = \frac{\partial L_{T-1}^B}{\partial \beta_2} + \frac{\partial L_T^B}{\partial \beta_2} > 0$$

Note that for when using the variable x^B the maximum likelihood will be found as parameter β_2 goes towards infinity, which is not the case for the x^A .

Now we proceed to the situation where the transition equation consists of both variables:

$$\pi_t = f(\alpha - \beta_1 x_t^A - \beta_2 x_t^B) \equiv f(\Theta)$$

What we will show is that the effect stemming from x_T^B will act very much more strongly than the effect of x_T^A . First, we consider the derivatives of the likelihood function:

$$\frac{\partial L}{\partial \beta_1} = \frac{\partial L_{T-1}}{\partial \beta_1} + \frac{\partial L_T}{\partial \beta_1} \quad (19)$$

and

$$\frac{\partial L}{\partial \beta_2} = \frac{\partial L_{T-1}}{\partial \beta_2} + \frac{\partial L_T}{\partial \beta_2} = 0 + \frac{\partial L_T}{\partial \beta_2} > 0$$

Again, the effect of this will be a solution in which $\beta_2 \rightarrow \infty$ and a more ambiguous solution for β_1 .

$$\lim_{\beta_2 \rightarrow \infty} \frac{\partial L}{\partial \beta_1} = \lim_{\beta_2 \rightarrow \infty} \frac{\partial L_{T-1}}{\partial \beta_1} + \lim_{\beta_2 \rightarrow \infty} \frac{\partial L_T}{\partial \beta_1} \quad (20)$$

where

$$\lim_{\beta_2 \rightarrow \infty} \frac{\partial L_{T-1}}{\partial \beta_1} = \frac{-x_{T-1}^A \exp(\Theta)}{[1 + \exp(\Theta)]^2} \cdot (a - b) = -x_{T-1}^A \underbrace{\frac{1}{[1 + \exp(\Theta)]}}_{\rightarrow 1} \underbrace{\frac{\exp(\Theta)}{[1 + \exp(\Theta)]}}_{\rightarrow 0} \cdot (a - b) = 0$$

and

$$\lim_{\beta_2 \rightarrow \infty} \frac{\partial L_T}{\partial \beta_1} = \frac{-x_T^A \exp(\Theta)}{[1 + \exp(\Theta)]^2} \cdot (b - a) = -x_T^A \underbrace{\frac{1}{[1 + \exp(\Theta)]}}_{\rightarrow 1} \underbrace{\frac{\exp(\Theta)}{[1 + \exp(\Theta)]}}_{\rightarrow 0} \cdot (b - a) = 0$$

Using these results, we see that (20) will converge towards zero. When maximizing the likelihood function, we will see the impact of the variable x_t^A diminish as the more and more weight is put on x_t^B through the parameter β_2 . It follows that the standard error of β_1 will become very large.

To summarize, the above discussion has assumed that we have a variable that is a perfect short-run predictor of future state switches. From this it has been shown that any other variable, although better resembling the true data generating process of transition probabilities, will be crowded out and deemed non-significant in a joint estimation.

Simulation evidence

Let us illustrate this in an applied setting. Consider the following simple 2 state model:

$$\Delta y_t = \mu_{S_t} + \epsilon_{S,t} \quad (21)$$

where $\bar{\mu} < 0$, $\mu_1 > 0 > \mu_2$, $\epsilon_t \sim N(0, \sigma_{S_t}^2)$, $\sigma_1^2 = 0.5$ and $\sigma_2^2 = 1$. The transition matrix is

$$\mathbf{P}_t = \begin{bmatrix} f(\alpha_1) & 1 - f(\alpha_1) \\ 1 - f(\alpha_2 + \beta y_{t-1}) & f(\alpha_2 + \beta y_{t-1}) \end{bmatrix}$$

In words, this process would have a negative drift were it not for the lower probability boundary that reverts the process back into the positive mean state. In figure 1 panel (a), a simulated series with the parameters $\mu_1 = 0.5$, $\mu_2 = -0.3$, $\beta = 0.1$ and $\alpha_1 = \alpha_2 = 3.4761 \Rightarrow p_{11} = 0.97$ is plotted.

To proceed, we have constructed a binary indicator variable x_t that takes on the value 1 the time period prior to a switch to state 1 and is 0 otherwise. We have then estimated three different setups of transition equations of the model:

- (i) $f(\alpha_2 + \beta_1 y_{t-1})$
- (ii) $f(\alpha_2 + \beta_1 y_{t-1} + \beta_2 x_{t-1})$
- (iii) $f(\alpha_2 + \beta_2 x_{t-1})$

It is apparent from table 1 that the predicted effect of including the binary indicator variable exists in the simulated data, even if the form of the long-run variable is different than from the theoretical set-up. The addition, the difference in the likelihood value between case (ii) and (iii) is virtually zero and the standard error of the parameter β_1 is very high indicating that the binary indicator dominates the long-run indicator.

The correlation coefficient has been computed as: $\rho = \text{Corr} [\Delta f(\Theta); \Delta f(\hat{\Theta})]$

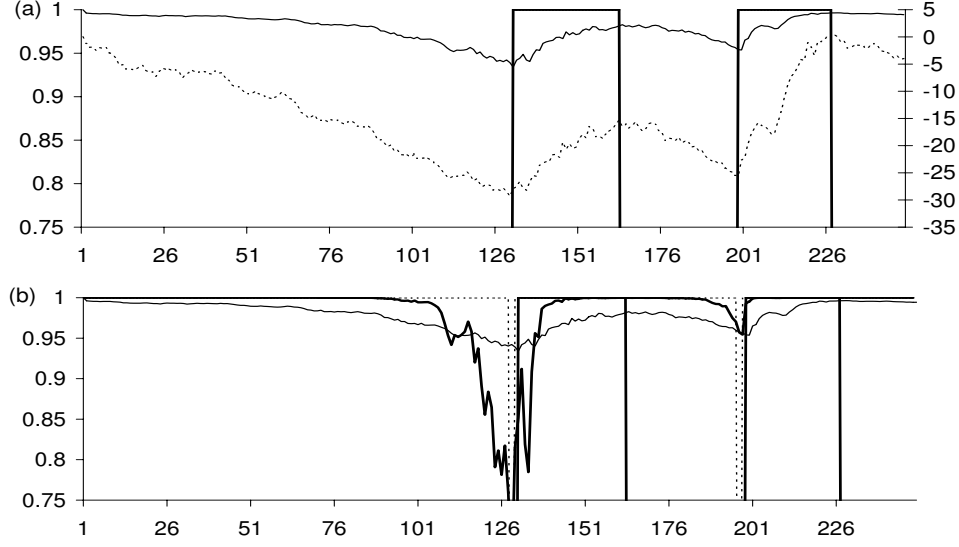


Figure 1: Simulated data. In panel (a): solid line indicates the y_t process (right scale); bars indicate the positive mean reversion state and the dotted line indicates the true transition probability process (left scale). In panel (b): Thin line marks the DGP transition probabilities as in panel (a), thick line marks the projection of transition probabilities from model (i) and dotted line marks the projection from model (iii). The projected transition probabilities from model (ii) are virtually identical to those of model (iii).

where Δ denotes the first difference operator and $\hat{\cdot}$ denotes empirical estimates. It indicates that the model without the short run indicator has a high correlation with the true process, whereas the other models - which should be preferred in terms of statistical significance - have much lower correlation coefficients. The importance of this effect can be seen in figure 1 where the true and projected transition probabilities have been plotted. As can be expected, the model including x_t signals a 0 probability to stay in state 1 one period prior to the actual shift but signals a stay probability of 1 otherwise. Model (i) shows a transition probability pattern similar to the DGP, but is more extreme in its projections. Specifically, it projects a much lower probability for transition to the second state when the actual transitions are 'far' off in the future, but the probability increases more rapidly as the switch comes closer, and close to the switch the probability becomes much higher than it is in the data-generating process. This can be indicative of an overfitting process, where the estimation process actually tries to mimic the behavior in the binary variable case. We saw that the binary case lead to a higher log likelihood value, so it is of no surprise that the estimation procedure behaves this way. The magnitude of this

Model	(i)		(ii)		(iii)		CTP	
Param.	Value	p	Value	p	Value	p	Value	p
μ_1	0.5967	0.00	0.6040	0.00	0.6040	0.00	0.6009	0.00
μ_2	-0.2255	0.00	-0.2273	0.00	-0.2273	0.00	-0.2260	0.00
σ_1	1.1414	0.00	1.1373	0.00	1.1373	0.00	1.1418	0.00
σ_2	0.4838	0.00	0.4846	0.00	0.4846	0.00	0.4833	0.00
α_1	3.3733	0.00	3.4369	0.00	3.4369	0.00	3.2719	0.00
α_2	21.3593	0.09	25	n.a.	14.6708	0.00	4.3380	0.00
β_1	0.7064	0.13	0.4221	0.94				
β_2			-25	n.a.	-25	n.a.		
LogL	-236.5861		-232.0960		-232.0961		-240.7682	
ρ	0.6311		0.2335		0.2335		n.a.	
R^2								

Table 1: Parameter values when models (i), (ii) and (iii) are estimated on the simulated data. The optimization procedure has been constrained to not allow parameters in the transition equation to exceed 25 in absolute value. Correlation denotes the correlation coefficient between the simulated TVP series and the empirically projected series.

effect will be explored further below.

Proposed Remedy

The previous investigation has shown that there is a bias towards selecting variables that induce changes in the transition equation very close and abruptly to a regime switch. We have argued that in empirical work, one may have the opposite objective. Also, including non-stationary variables in the transition equation results in an inference problem similar to that of standard spurious regressions. In this section, we will suggest a simple solution to remedy these problem. To do this, we require the researcher's prior about how important the long-run effects are in relation to the short-run ones.

We begin with the cases considered initially, cases (i) and (ii). As we saw both in the theoretical and empirical setting, the problem with case (ii) was that maximization of the likelihood function led to a corner solution for β_2 . A natural way to avoid this is to introduce a penalty in the likelihood function so that there exist a finite solution for (18) = 0. The problem is to decide upon the magnitude and functional form of the penalty. The simple approach suggested here uses a prior about how the projected transition probability series should look. To begin with, we assume that fundamental economic variables evolve slowly over time. Then, if these fundamentals govern the probability to switch economic states, we would expect the series of probabilities to correspondingly move slowly. In a graphical depiction of the probabilities, a smooth series implies slow movements in the under-

lying variable one measures. For example, the Hodrick-Prescott filter decomposes a time-series into a slowly moving, smooth trend component and a faster moving, non-smooth cyclical component.

Hence, we proxy the prior of slowly moving fundamentals with a term in the likelihood function that penalizes non-smooth behavior. The penalized log likelihood takes on the following form:

$$L(\theta) = \sum_{t=1}^T \left\{ \log \mathbf{1}' \left(\xi_{t|t-1} \odot \eta_t \right) - e^\gamma \mathbf{1}' [\text{diag}(\mathbf{P}_t) - \text{diag}(\mathbf{P}_{t-1})]^2 \right\} \quad (22)$$

where $\text{diag}(\cdot)$ denotes the principal diagonal operator and γ is the weight of the prior given by the econometrician. The first drawback of this approach is obvious: traditional likelihood ratio testing will not be possible using the expression in (22) since the penalty term will make it be lower than the baseline model's log likelihood. However, since this change of the likelihood is always negative, a likelihood ratio statistic based on it will be more conservative in the sense that it rejects too many variables. One way to reduce this drawback is to use the estimation results obtained from maximizing (22) and evaluate the non-penalized likelihood function with the corresponding parameter vector.⁴ Still, the size of the test will be biased downwards. Erlandsson (2004) shows, however, that even in relatively large samples, the likelihood ratio test for the transition equation regressors is oversized in the time-varying transition probability Markov switching model. Consequently, the penalty can actually have the effect of bringing the test closer to its nominal size, than otherwise. The exact magnitudes of the downward and upward biases remain to be investigated, but it should be noted that the upward bias on the size is positively correlated to sample size and persistence of the transition regressor *ceteris paribus*.

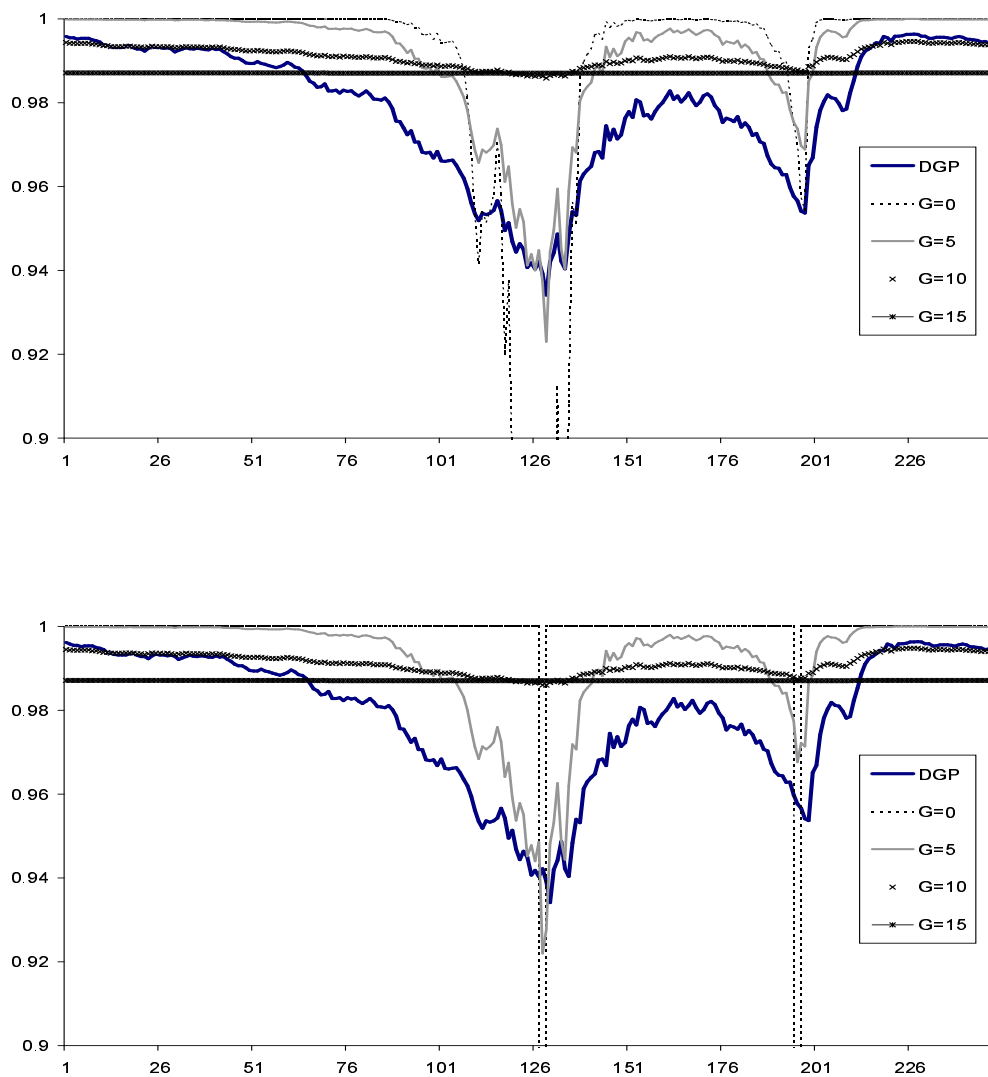
Both panels of figure 2 show how applying the penalty results in a more smooth transition probability function. It also shows the trade-off between smoothness and magnitude of the predicted probabilities. Once the functions become more smooth, it is less capable of inducing a large transition probability. For penalties of 15 and more, the stay probability is always more than 98%. In table 2, the results show - up to a certain level for the penalty - that the penalized models increase the correlation between the true DGP and the projected stay probabilities. The coefficient estimates are reduced and come closer to their true value as well, and the modified likelihood values decrease as the penalty increases.

From the bottom panel of figure 2, where model (ii) has been estimated and used to project stay probabilities, the importance of the penalty becomes more protruding. The unconstrained model has the binary looking transition series, whereas the prior constrained series exhibit patterns (and by looking in table 3 correlations) closely linked to the true DGP.

So far, we have only seen the effect of the penalty in one data-set. The next logical step is to study the effects in a generalized setting. We now use a new

⁴The latter case will be denoted with an additional *.

Figure 2: Stay probabilities with different priors ($\gamma = G$) ; model (i) (upper panel) and model (ii) (bottom panel).



	α_2	β_1	ρ^Δ	R^2	MSE	Log-L*
$\gamma = 0$	19.8316	0.6570	0.7143			-235.7250
$\gamma = 1$	16.7850	0.5375	0.7579			-235.7753
$\gamma = 3$	13.0636	0.3874	0.8323			-236.0339
$\gamma = 5$	9.9710	0.2556	0.9184			-236.5826
$\gamma = 7$	7.8733	0.1574	0.9826			-237.3949
$\gamma = 10$	5.2109	0.0330	0.9739			-239.0265
$\gamma = 15$	4.3393	0.0002	0.9403			-239.7009

Table 2: Diagnostics for different penalty priors; model (i), simulated data set as in figure 1.

	α_2	β_1	β_2	ρ^Δ	Log-L*
$\gamma = 0$	23.8403	0.3839	-25.0000	0.1811	-231.0568
$\gamma = 1$	25.0000	0.0000	-24.1682	0.1811	-233.4431
$\gamma = 3$	12.9989	0.3768	-0.6987	0.8039	-235.6277
$\gamma = 5$	10.1913	0.2608	-0.3066	0.9064	-236.3956
$\gamma = 7$	8.0139	0.1612	-0.1501	0.9784	-237.2978
$\gamma = 10$	5.2474	0.0339	-0.0320	0.9751	-238.9969
$\gamma = 15$	4.3394	0.0002	-0.0002	0.9410	-239.7007

Table 3: Diagnostics for different penalty priors; model (ii), simulated data set as in figure 1.

set of parameters for the model in (21). We set $\mu_1 = 1$, $\mu_2 = -1$, $\sigma^2 = 1$, $\alpha_1 = 3.89 \Leftrightarrow p_{11} = 0.98$, $\alpha_2 = 5$ and $\beta = -0.1$. We simulate a large number of data-series based on these parameters and with sample sizes 150, 250, 500 and 1000. We then estimate the model with penalties 0,1,3,5,7 and 10 and produce a set of statistics for the transition probability projections vis-a-vis the true transition probabilities.⁵ The results are collected in figure 3.

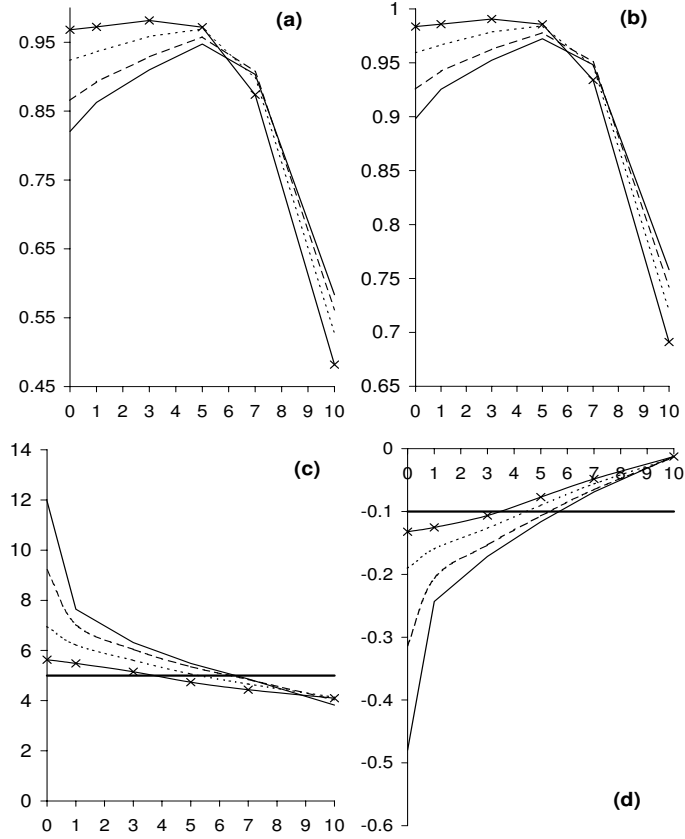


Figure 3: Effects of penalty. Sample sizes 150 (solid), 250 (long broken), 500 (broken) and 1000 (crosses). Panel (a) outlines the average R^2 values in a regression of the TVP projection and a constant on the data-generating transition probabilities. Panel (b) contains the average ρ^Δ measure; panels (c) and (d) the average estimates of α_2 and β .

The tendencies outlined in the simple example are clearly systematic. First, we

⁵We drop the penalty $\gamma = 15$ since those estimations often run into numerical problems.

note that sample size has a marked effect on the impact of the penalty. For the 150 observation sample, the R^2 of the TVP regression rises from .821 in the non-penalized model to .947 when $\gamma = 5$. The R^2 for the longest sample rises from .968 ($\gamma = 0$) to .982 ($\gamma = 3$), which is a much lower increase.

Another result with implications for selecting γ protrudes from panels (c) and (d). We note that for the shortest sample length, $\gamma \approx 6$ yields coefficient estimates close to the true ones. As sample-length increases the necessary γ for this condition to be fulfilled decreases. For the 1000 observation sample length, $\gamma \approx 3$ seems to be more appropriate. When the R^2 and correlation graphs are more closely scrutinized, one can see this pattern also appearing there. Again, this should be related to the size-distortions found in Erlandsson (2004): a larger penalty means a larger downward bias in the size of the test. But, inversely, a shorter sample leads to a higher upward size bias.

In the end, this could indicate a future possible way to pick an "optimal" value for γ : one could simply choose the γ that reduces the size-distortion of the likelihood ratio test. Whether this would actually generate the γ that maximizes the correlation between projected transition probabilities and the true ones remains to be studied. This may also conflict with the objective of obtaining long-run predictors. Hence, for the remainder of the paper we will hold on to the subjective selection of γ .

Empirical Application

The original Hamilton (1989) paper established the usefulness of the Markov regime switching model to replicate business cycles. A wider discussion of the relevance of Markov switching in modelling asymmetrical GDP growth is available in Hamilton and Raj (2002). We choose to study quarterly real GDP data from 1964:1 to 2002:4 for a total of 150 observations.⁶, which is a similar data-set to that studied in Coe (2002). The baseline model is

$$\Delta y_t = \mu_{S_t} + \epsilon_{R_t} \quad (23)$$

where Δy_t is the logarithmic change in real GDP per capita and S_t, R_t are unobserved state variables. The error term ϵ_{R_t} is distributed according to $N(0, \sigma_{R_t}^2)$. The first state variable, $S_t \in [S^{Expansion}, S^{Contraction}]$; $S^{Contraction} \prec S^{Expansion}$, governs the intercept and the second one R_t governs volatility. Initially, the transition matrix \mathbf{P}_t is kept constant as in equation (2). The transition matrix \mathbf{Q} associated with the R process is assumed to be constant through out the remainder of the paper.

In order to test for the existence of Markov switching dynamics in the data, we apply the Monte Carlo testing procedure discussed in Cheung and Erlandsson

⁶Nominal GDP and inflation as measured by the consumer price index are obtained through the IMF's *International Financial Statistics* database.

	LR	H_0^N	H_0^M
$N = 1; M = 2$	33.9970	0.0040	0.5737
$N = 2; M = 4^*$	15.8450	0.0080	0.9920

Table 4: Test for Markov switching dynamics. 250 Monte Carlo runs.

(2005), which is an extension of the Rydén, Teräsvirta and Åsbrink (1998) procedure. The results in table 4 indicate strong evidence of Markov switching dynamics in the data. Diagnostic testing rejects the hypothesis that variances are equal across states for the standard 2 state setting. Consequently, we also investigate the possibility of variance following a regime switching process of its own so that $S_t \neq R_t$ for some t , thus allowing for a total of 4 states.⁷ The results of this test are also clear. We reject the 2 state MS model, but are unable to reject the 4* state counterpart. One could also suspect even more states in the data, but limited computational capacity restricts us from investigating these suspicions. An alternative is to look at the diagnostics of the model in the proposed specification. With 4* states, neither significant residual autocorrelation as measured by the Ljung-Box Q statistic (p-value 0.074, nor ARCH effects as measured by Engle’s LM test (p-value 0.173), is present in the standardized residuals. For the 2 state model the corresponding p-values are 0.001 and 0.311 respectively. Thus, it seems unnecessary to add more states in order to capture the dynamics of the first, second and fourth moments.

Another diagnostic measure to validate the model is how well it replicates business cycles as measured elsewhere. The by all standards most common benchmark in the literature is the National Bureau of Economic Research (NBER) business cycle dates., which we will denote as \hat{S}_t . The smoothed probabilities, computed according to the algorithm of Kim (1994) of the the contractionary state of real-GDP in our model is plotted against the NBER dates in figure 4. As can be seen, the model replicates the dates quite well. Only using 2 states produces a graph that does not resemble the NBER dates.⁸ The reason for this seems to be a shift from the low volatility state to the high volatility states in 1984:1-1984:2. The high volatility state seems to be absorbant within the sample, meaning that the volatility process does never return to the low state after 1984. The simpler model produces probabilities that are a mix of the level and volatility states in the more general model.

To proceed with investigating factors that predict recession, we convert the baseline model to a restricted TVP parameterization. The transition matrix for the S_t process is

$$\mathbf{P}_t = \begin{bmatrix} f(\alpha_1 + \beta \mathbf{X}_t) & 1 - f(\alpha_1 + \beta \mathbf{X}_{t-1}) \\ 1 - f(\alpha_2) & f(\alpha_2) \end{bmatrix}$$

⁷This parameterization is constricted however so that the 4 state transition matrix is the Kronecker product of the two separate processes’ respective transition matrices. Since the resulting transition matrix is constrained, we will denote these 4 states with a subscript $* \rightarrow 4^*$.

⁸Graph available upon request.

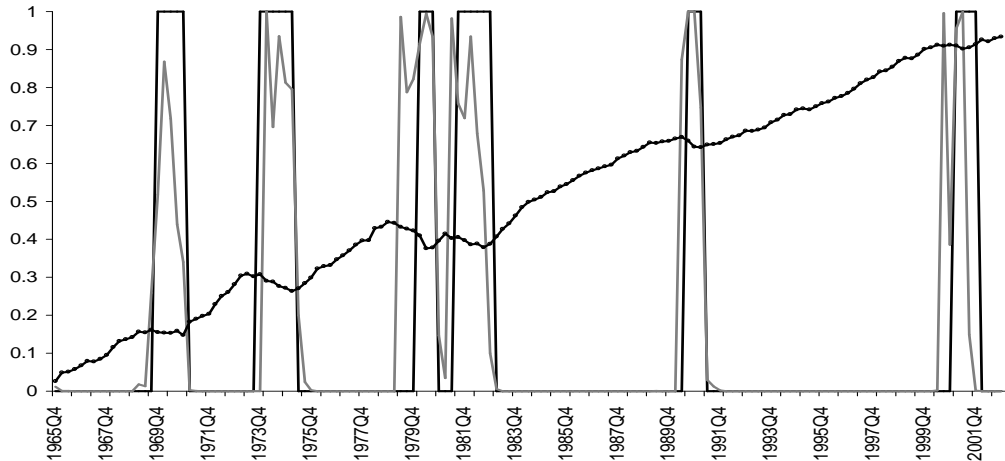


Figure 4: NBER dates (black) and smoothed probabilities of the contractionary state (grey) using the constant transition probability model. U.S. real GDP is plotted with dots (normalized to 0 in 1965Q4).

where \mathbf{X}_{t-1} is a set of possible leading indicator candidates. This structure means that the transition probability from the expansionary phase of the economy to the contractionary is time varying, whereas the reverse is constant.

The next step is to specify \mathbf{X} . In table 8, we present 31 variables suggested by *Econmagic* to be related to the business cycle. Monthly data has been transformed to quarterly by taking the end of quarter monthly value. Each variable in first differences, an 8 quarter moving average, and in levels has been tested individually through likelihood ratio tests,⁹ and with differing penalty terms. 12 of the candidates have median p-values below 20%.¹⁰ Of these, 10 are 8 quarter moving averages and 2 are in levels. This should be viewed in the light that 18 variables in first differences are significant at the 10% level when estimating the model without a penalty term, but all of them turn insignificant once the penalty term is applied.

We choose to be conservative when setting γ . A low γ will mean smaller deviances from the standard application procedure, but given the results in our earlier simulations we should be able to outline the positive effects of the penalty even at low levels of γ . Consequently, we set $\gamma = 2$. Using this prior, we have conducted

⁹To obtain better convergence properties all variables have been normalized. The level variable has been calculated as the cumulative sum of normalized first differences so that any time trends have been removed. The moving average has been calculated the same way, but as the average of the 8 last observations.

¹⁰The median p-values are calculated as the median of the p values for one variable, one transformation and 6 different penalty settings (ranging from 0 to 5).

a testing down procedure of \mathbf{X} . The least significant variable has been removed until the reduction in the likelihood ratio statistic is below 5% level.¹¹ The final specification that are reached are presented in table 5, column 2. The results from a number of benchmark models are also presented. The first column, $\gamma = 0$, refers to the results when estimating the specification obtained from the penalized setting but setting $\gamma = 0$. The CTP model is the constant transition probability model. In the fourth column, denoted $\gamma = 0(*)$, a final specification has been obtained using the same testing down procedure as above, but with $\gamma = 0$.

The results indicate that all TVP model estimates have more severe contractionary phases than in the CTP case. When using the penalized version, two variables are found significant as leading indicators of a recession: the seasonally adjusted production price index of finished goods and the number of unemployed civilians.¹² A rise in the production prices decreases the probability to stay in the expansionary state, as does an increase in unemployment. One interpretation of these indicators is that the probability of recessions is strongly linked to shortages in both goods and labor markets.

Using the traditional approach of testing down the model's TVP variables, a very different conclusion is reached. First, many more variables are deemed significant, which also was the prediction of the simulation results in the previous section. Second, for the one variable that the specification have in common, the signs are opposite. The difference in likelihood values is large, even when applying the non-penalized value on the penalized specification.

Using the above specifications of the model and its benchmarks, we register dates t at which the predicted probability to proceed to the contraction state in the time interval $t + 1 : t + k$ exceeds a certain threshold level in percent, denoted ω . The forecasted probability for $t + 1$ is calculated as:

$$\Pr(S_{t+1} = 1|\Omega_t) = p_{11,t}\Pr(S_t = 1|\Omega_t) + p_{21,t}\Pr(S_t = 2|\Omega_t) \quad (24)$$

For the $k > 1$ step ahead forecast, we focus on the probability that we will at least one crisis period within the time interval $t + 1, t + 2, \dots, t + k - 1, t + k$. This equals 1 minus the probability that we see no crises within the time interval:

$$\begin{aligned} \Pr(\min(S_{t+1}, \dots, S_{t+k}) = 1|\Omega_t) &= 1 - \Pr(S_{t+1}, \dots, S_{t+k} = 2|\Omega_t) = \\ &= 1 - \left[p_{12,t} p_{22,t}^{k-1} \Pr(S_t = 1|\Omega_t) + p_{22,t}^k \Pr(S_t = 2|\Omega_t) \right] \end{aligned} \quad (25)$$

Four cases of signals from the model vis-a-vis the actual development can then be constructed:¹³

¹¹In this non-linear setting, the likelihood ratio statistic has been shown to be more robust than statistics based on the variance-covariance matrix, such as the Wald statistic.

¹²The latter variable reflects the number of unemployed civilians compared to the trend, and is consequently similar to an ordinary unemployment rate figure.

¹³The first letter in each case's acronym stands for the prediction of the model, the second for the actual development. C refers to at least one contraction within the time interval, E (as in expansions) to a time interval with no contraction.

	$\gamma = 0$		$\gamma = 2$		CTP		$\gamma = 0(*)$	
	Value	p	Value	p	Value	p	Value	p
$\mu^{Expansion}$	0.7308	0.00	0.7394	0.00	0.7459	0.00	0.7396	0.00
$\mu^{Contraction}$	-0.7681	0.00	-0.7118	0.07	-0.6064	0.14	-0.9062	0.00
σ_1^2	1.0213	0.00	1.0186	0.00	1.0397	0.00	0.9701	0.00
σ_2^2	0.4866	0.00	0.4818	0.00	0.4765	0.00	0.4759	0.00
α_1^S	100	0.00	3.8072	0.00	2.9896	0.00	100	n.a.
β_{13}^S							-1.6376	0.00
β_{17}^S							0.6659	0.00
β_{20}^S							-0.4281	0.00
β_{21}^S							-0.5264	0.00
β_{22}^S	-0.743	0.00	-0.0129	0.00				
β_{23}^S							1.4974	0.00
β_{25}^S	0.9836	0.01	0.0157	0.00			-0.2270	0.00
β_{26}^S							-0.5538	0.00
α_2^S	0.2332	0.80	0.7867	0.26	0.9697	0.13	-0.6939	0.68
α_1^R	4.2969	0.00	4.2967	0.00	4.2960	0.00	4.2981	0.00
α_2^R	100	n.a.	100	n.a.	100	0	100	n.a.
LogL	-167.84		-175.77		-181.06		-161.52	
R ²	0.5775		0.5825		0.5494		0.5973	
LR	26.4300	0.00	10.5670	0.01	0	n.a.	39.0743	0.00
LR*	26.4300	0.00	12.9619	0.00	0	n.a.	39.0743	0.00

Table 5: Estimation results. The optimization procedure has been bounded so that $-100 \leq \alpha, \beta \leq 100$. Standard errors of the stay probability parameter in the low volatility state, α_2^R , are not computable since the state is absorbant within the data range. Subscript indices on β refers to the index number of exogenous variables in table 8.

1. The model signals a contraction, and a contraction occurs (CE):
 $\Pr(\min(S_{t+1,\dots,t+k}) = S^{Contraction} | \Omega_t) > \omega \text{ and } \min(S_{t+1:t+k}) = S^{Contraction}$
2. The model signals a contraction, but no contraction occurs (CE):
 $\Pr(\min(S_{t+1,\dots,t+k}) = S^{Contraction} | \Omega_t) > \omega \text{ and } \min(S_{t+1:t+k}) = S^{Expansion}$
3. The model signals no contraction, but a contraction occurs (EC):
 $\Pr(\min(S_{t+1,\dots,t+k}) = S^{Contraction} | \Omega_t) \leq \omega \text{ and } \min(S_{t+1:t+k}) = S^{Contraction}$
4. The model signals no contraction, and no contraction occurs (EE):
 $\Pr(\min(S_{t+1,\dots,t+k}) = S^{Contraction} | \Omega_t) \leq \omega \text{ and } \min(S_{t+1:t+k}) = S^{Expansion}$

Using these definitions, a number of benchmarks of the model's performance can be constructed. We will focus on two; the first being the ratio of correct signals to the total number of signals $(CC + EE)/(CC + CE + EC + EE)$. In our setting this benchmark answers the question: "How reliable is the predictions of the model?" The second benchmark, the noise-to-signal ratio, is calculated as $[EC/(EC + EE)]/[CC/(CC + CE)]$, and returns a measure of how strong inference on the true development the model gives. A model that perfectly predicts the future would have a noise-to-signal ratio of 0. The threshold for when a signal is given, ω , may be set to different values than the traditional 50% to reflect different sensitivities to recessions.

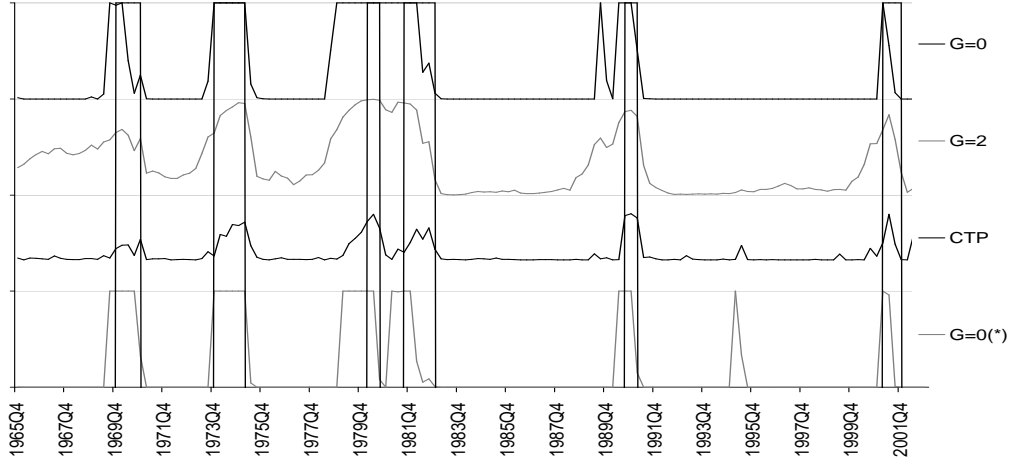


Figure 5: Forecasts of the probability that a contraction will occur within the next 8 quarters.

Table 6 presents the results of various models. We first note that the only instance where the CTP model outperforms the TVP model in terms of correct forecasts is for the 1 quarter ahead predictions with a threshold of 50%. For prediction horizons of greater than 4 quarters, the CTP model exhibits much worse performance than the TVP models. This is a matter of pure arithmetics: with the 50/25% threshold, the model always predicts the probability to enter the contractionary phase within 12/8 quarters to be greater than the threshold.

The trend for the TVP models based on variable selection found with the penalized likelihood function (i.e. the cases $\gamma = 0$ and $\gamma = 2$) is that as the prediction horizon expands, the better the $\gamma = 2$ model is relative to the $\gamma = 0$ model. This also hold for the relation between the $\gamma = 2$ and $\gamma = 0(*)$ models. Looking at what types of errors the models make, we note that the penalized model for these horizons predicts a larger number of recessions, resulting in fewer case 3 errors.

Figure 5 provides a graphical illustrations on the forecasting performance at the 8 quarter horizon. As has been sought for, the penalized model exhibits a much smoother projection of transition probabilities than the benchmark models. The probabilities also seem to rise earlier prior to recession than for the benchmarks.

A more reliable way to evaluate the model's performance is to observe the out-of-sample forecasting properties. If the model reflects a relation that is stable over time, out-of-sample forecasts will resemble the corresponding in-sample forecasts. Otherwise, we have an indication of overfitting. The out-of-sample forecasting performance is tabulates in table 7. We have selected to produce forecasts of the two last recessions in the sample, ending the in-sample at 1988:2 and forecasting 1988:3

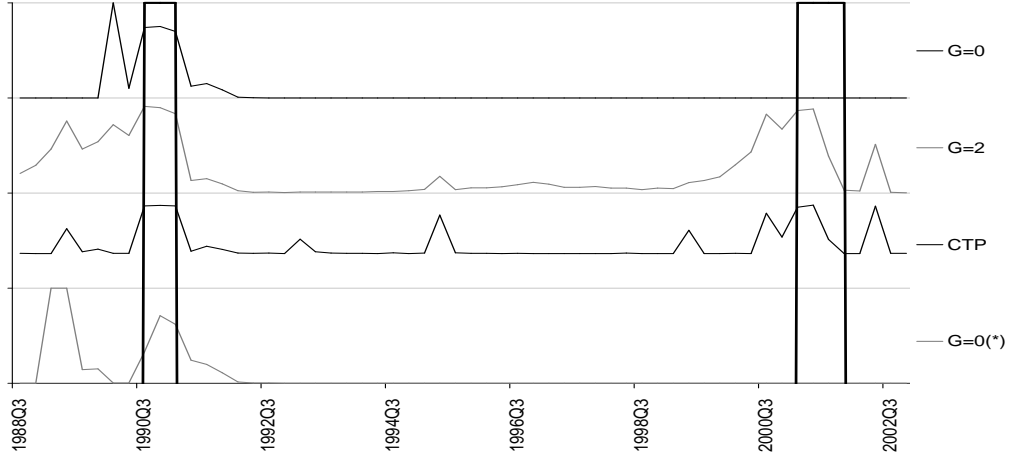


Figure 6: Out-of-sample forecasts of the probability that a contraction will occur within the next 8 quarters.

to 2002:4. Since the latest recession has occurred so recently, this means that the performance of the 8 quarter and 12 quarter forecasts cannot be evaluated for that recession.

The penalized model consistently outperforms the benchmark models in this setting at all horizons. When viewing the graphical evidence as in figure 6, it turns out that the penalized model is alone among the time varying probability models in being able to predict the 2001 recession. It is also more consequent in predicting the 1990 recession, with a gradual increase in probabilities rather than the jagged projection of the other TVP models.

Conclusion

This paper illustrates the inability of the Markov regime switching model to make inference on the probability of states occurring in the medium to long term. For policy purposes, short run predictors of future states may be irrelevant, since many tools such as fiscal policy and interest rate changes only have effects in the medium to long term. Hence, the model has not been a commonly used tool when predicting future states of the economy.

Rather than not being able to make inference in the longer run at all, the model with time varying transition probabilities possesses a bias towards selecting short run variables for predicting future states. This also leads to estimation problems in the maximum likelihood setting, where bounded optimization procedures often has

parameter estimates of the TVP variables on the boundaries. Summed together: the estimated model seems to be unable to depict a highly useful dimension of the theoretical model, and it may be hard to obtain estimates at all.

We propose a simple penalty term, based on the smoothness of the projected time varying probabilities, in the maximum likelihood function which aims to remedy this problems. Simulation evidence indicates the usefulness of the penalty: the correlation between the true transition probability process, and the projections obtained from the penalized model is notably higher than for the non-penalized model. This effect is diminishing as the sample size grows. We also argue that the penalty's distortion of the distributional properties of the likelihood ratios may very well just counterweigh the distortion of the test shown elsewhere to appear in limited samples.

In an empirical application, we use a number of suggested leading indicators to predict contractionary states of U.S. real GDP. The standard ML estimates are shown to possess the problems shown in the simulation exercise. A large number of variables show up as significant, and the projected transition probabilities are very non-smooth. Applying the proposed penalized estimator yields a final model specification with fewer variables and smooth transition probabilities. In the in-sample forecasting exercise, the penalized model performs better for longer (8-12 quarter) horizons. When calculating out-of-sample forecasts, the penalized models exhibits better performance irrespective of the horizon. It is the only model that is able to predict the 2001 recession out-of-sample.

References

- Abiad, A. (2003). Early-Warning Systems: A Survey and a Regime-Switching Approach. IMF Working Paper 03/32.
- Cheung, Y.-W. and Erlandsson, U. G. (2004). Exchange rates and Markov switching dynamics. Forthcoming in Journal of Business and Economic Statistics.
- Coe, P. (2002). Power Issues When Testing the Markov Switching Model with the Sup Likelihood Ratio Test Using U.S. Output. *Empirical Economics*, 27:395–401.
- Diebold, F. X., Lee, J.-H., and Weinbach, G. (1994). Regime switching with time-varying transition probabilities. In Hargreaves, C., editor, *Nonstationary Time Series Analysis and Cointegration*, pages 283–302, Oxford. Oxford University Press.
- Erlandsson, U. G. (2004). Transition Variables in the Markov Switching Model: Some Small Sample Properties. Manuscript, Department of Economics, Lund University.
- Filardo, A. J., , and Gordon, S. F. (1998). Business cycle durations. *Journal of Econometrics*, 85:99–123.

- Gray, S. F. (1996). Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42:27–62.
- Hamilton, J. and Raj, B. (2002). New directions in business cycle research and financial analysis. *Empirical Economics*, 27:149–162.
- Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57:357–384.
- Hamilton, J. D. (1990). Analysis of Time Series Subject to Changes in Regime. *Journal of Econometrics*, 45:39–70.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton.
- Karlsen, H. (1990). *A Class of Non-Linear Time Series Models*. PhD thesis, Department of Mathematics, University of Bergen, Norway.
- Kim, C.-J. (1994). Dynamic Linear Models with Markov-Switching. *Journal of Econometrics*, 60:1–22.
- Rydén, T., Teräsvirta, T., and Åsbrink, S. (1998). Stylized facts of daily return series and the hidden Markov model. *Journal of Applied Econometrics*, 13:217–244.
- Schaller, H. and van Norden, S. (2002). Fads or bubbles? *Empirical Economics*, 27:335–362.
- Tronzano, M., Psaradakis, Z., and Sola, M. (2003). Target zone credibility and economic fundamentals. *Economic Modelling*, 20:791–807.

Threshold 50%	$\gamma = 0$	$\gamma = 2$	CTP	$\gamma = 0(*)$
$k = 1$				
Correct obs. ratio	0.8919	0.8716	0.9054	0.8581
Noise-to signal ratio	0.0853	0.1491	0.0856	0.1776
# CC/CE/EC/EE	20/9/7/112	13/5/14/116	19/6/8/115	10/4/17/117
$k = 4$				
Correct obs. ratio	0.8621	0.8414	0.8414	0.8483
Noise-to signal ratio	0.1667	0.1912	0.1915	0.1831
# CC/CE/EC/EE	27/2/18/98	24/2/21/98	25/3/20/97	25/2/20/98
$k = 8$				
Correct obs. ratio	0.7376	0.7943	0.6809	0.7092
Noise-to signal ratio	0.3356	0.2770	0.4612	0.3654
# CC/CE/EC/EE	27/2/35/77	37/4/25/75	27/10/35/69	23/2/39/77
$k = 12$				
Correct obs. ratio	0.6131	0.7445	0.5693	0.5839
Noise-to signal ratio	0.5072	0.3943	n.a.	0.5338
# CC/CE/EC/EE	27/2/51/57	48/5/30/54	78/59/0/0	22/2/55/57
Threshold 25%	$\gamma = 0$	$\gamma = 2$	CTP	$\gamma = 0(*)$
$k = 1$				
Correct obs. ratio	0.9054	0.8919	0.8716	0.9122
Noise-to signal ratio	0.0499	0.0853	0.0601	0.0578
# CC/CE/EC/EE	23/10/4/111	20/9/7/112	23/15/4/106	22/8/5/113
$k = 4$				
Correct obs. ratio	0.8690	0.8897	0.7103	0.8414
Noise-to signal ratio	0.1532	0.0562	0.2075	0.1906
# CC/CE/EC/EE	30/4/15/96	41/12/4/88	37/34/8/66	26/4/19/96
$k = 8$				
Correct obs. ratio	0.7589	0.8227	0.4397	0.7021
Noise-to signal ratio	0.3131	0.1938	n.a.	0.3923
# CC/CE/EC/EE	30/2/32/77	50/13/12/66	62/79/0/0	24/4/38/75
$k = 12$				
Correct obs. ratio	0.6423	0.7956	0.5693	0.5912
Noise-to signal ratio	0.4811	0.3031	n.a.	0.5446
# CC/CE/EC/EE	31/2/47/57	62/12/16/47	78/59/0/0	25/3/53/56

Table 6: In-sample prediction results.

Threshold 50%	$\gamma = 0$	$\gamma = 2$	CTP	$\gamma = 0(*)$
$k = 1$				
Correct obs. ratio	0.8772	0.8772	0.8772	0.8421
Noise-to signal ratio	0.1887	0.1224	0.1224	0.4528
# CC/CE/EC/EE	2/2/5/48	4/4/3/46	4/4/3/46	1/3/6/47
$k = 4$				
Correct obs. ratio	0.7963	0.8333	0.7778	0.7222
Noise-to signal ratio	0.2667	0.2029	0.3200	0.9600
# CC/CE/EC/EE	3/1/10/40	6/2/7/39	5/4/8/37	1/3/12/38
$k = 8$				
Correct obs. ratio	0.6800	0.7800	0.7000	0.6800
Noise-to signal ratio	0.4348	0.2744	0.4390	0.4348
# CC/CE/EC/EE	3/1/15/31	8/1/10/31	6/3/12/29	3/1/15/31
$k = 12$				
Correct obs. ratio	0.6522	0.7609	0.5870	0.6522
Noise-to signal ratio	0.4762	0.3041	0.7879	0.4762
# CC/CE/EC/EE	3/1/15/27	8/1/10/27	6/7/12/21	3/1/15/27
Threshold 25%	$\gamma = 0$	$\gamma = 2$	CTP	$\gamma = 0(*)$
$k = 1$				
Correct obs. ratio	0.8772	0.8947	0.8596	0.8596
Noise-to signal ratio	0.1887	0.0750	0.0492	0.2404
# CC/CE/EC/EE	2/2/5/48	5/4/2/46	6/7/1/43	2/3/5/47
$k = 4$				
Correct obs. ratio	0.7963	0.9074	0.7778	0.7407
Noise-to signal ratio	0.2667	0.0636	0.2404	0.5612
# CC/CE/EC/EE	3/1/10/40	11/3/2/38	8/7/5/34	2/3/11/38
$k = 8$				
Correct obs. ratio	0.6800	0.8800	0.3600	0.7000
Noise-to signal ratio	0.4348	0.1496	n.a.	0.3889
# CC/CE/EC/EE	3/1/15/31	13/1/5/31	18/32/0/0	4/1/14/31
$k = 12$				
Correct obs. ratio	0.6522	0.8043	0.3913	0.6739
Noise-to signal ratio	0.4762	0.2514	n.a.	0.4268
# CC/CE/EC/EE	3/1/15/27	10/1/8/27	18/28/0/0	4/1/14/27

Table 7: Out-of-sample prediction results.

Indicator	Reporting frequency	Index number
Total private: Indexes of Aggregate Weekly Hours, SA	M	1
Average Weekly Hours; Private Nonagricultural Establishments; SA	M	2
Total Borrowings at Federal Reserve Banks; Billions of Dollars; NSA	M	3
Change in Business Inventories; SAAR Billions of Dollars	Q	4
Corporate Profits After Tax with IVA and CCAdj; Billions; SAAR	Q	5
Consumer Price Index All Urban Consumers: Total; 1982-84=100; SA	Q	6
Consumer Price Index All Urban Consumers: Less Food and Energy; 1982-84=100, SA	M	7
Employment Ratio; Civilian Employment/Civilian Non. Inst. Pop.; Percent SA	M	8
Gross Savings; Billions of Dollars SAAR	Q	9
Index of Help Wanted Advertising; in Newspapers; 1987=100; SA	M	10
Total Industrial Production Index; 1992=100 SA	M	11
M2 Money Stock; Billions of Dollars; SA	M	12
Bank Prime Loan Rate	M	13
Nonfarm Business Sector: Output Per Hour of All Persons; SA, 1992=100	Q	14
Private Business Sector: Output Per Hour of All Persons; SA, 1992=100	Q	15
Payroll Employment; of Wage and Salary Workers; Thousands; SA	M	16
Personal Consumption Expenditures; Billions of Dollars SAAR	M	17
Personal Income; Billions of Dollars SAAR	M	18
PPI - Capital Equipment; 1982=100 SA	M	19
PPI - Crude Materials for Further Processing; 1982=100 SA	M	20
PPI - Finished Consumer Foods; 1982=100 SA	M	21
PPI - Finished Goods; 1982=100 SA	M	22
PPI - Intermediate Materials; 1982=100 SA	M	23
Personal Saving; Billions of Dollars SAAR	Q	24
Civilian Unemployed for 15 Weeks and Over; Thousands; SA	M	25
Manufacturing Sector: Unit Labor Cost; SA, 1992=100	Q	26
Nonfarm Business Sector: Unit Labor Cost; SA, 1992=100	Q	27
Consumer Sentiment; University of Michigan; 1966Q1=100; NSA	Q	28
Unemployment Level; All Civilian Workers; Thousands; SA	Q	29
Capacity Utilization: Manufacturing (SIC); SA	M	30
Industrial Production Index: Consumer goods; 1997=100; SA	M	31

Table 8: Evaluated predictors of the transition probability to the contraction state.