CIRANO
Centre interuniversitaire de recherche
en analyse des organisations

Série Scientifique
Scientific Series

N° 94s-15

BAYESIAN INFERENCE FOR PERIODIC REGIME-SWITCHING MODELS

Eric Ghysels, Robert E. McCulloch, Ruey S. Tsay

Montréal
Novembre 1994

CIRANO

Le CIRANO est une corporation privée à but non lucratif constituée en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de l'Industrie, du Commerce, de la Science et de la Technologie, de même que des subventions et mandats obtenus par ses équipes de recherche. La *Série Scientifique* est la réalisation d'une des missions que s'est données le CIRANO, soit de développer l'analyse scientifique des organisations et des comportements stratégiques.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère de l'Industrie, du Commerce, de la Science et de la Technologie, and grants and research mandates obtained by its research teams. The Scientific Series *fulfils one of the missions of CIRANO: to develop the scientific analysis of organizations and strategic behaviour.*

Les organisations-partenaires / The Partner Organizations

- Ministère de l'Industrie, du Commerce, de la Science et de la Technologie.
- École des Hautes Études Commerciales.
- École Polytechnique.
- Université de Montréal.
- Université Laval.
- McGill University.
- Université du Québec à Montréal.
- Bell Québec.
- Caisse de dépôt et de placement du Québec.
- Hydro-Québec.
- La Banque Laurentienne du Canada.
- Fédération des caisses populaires de Montréal et de l'Ouest-du-Québec.

ISSN 1198-8177

Bayesian Inference for Periodic Regime-Switching Models

Eric Ghysels[†]
Robert E. McCulloch[‡]
and
Ruey S. Tsay[‡]

Abstract / Résumé

We present a general class of nonlinear time series Markov regime-switching models for seasonal data which may exhibit periodic features in the hidden Markov process as well as in the laws of motion in each of the regimes. This class of models allows for nontrivial dependencies between seasonal, cyclical and long-term patterns in the data. To overcome the competitional burden we adopt a Bayesian approach to estimation and inference. This paper contains two empirical examples as illustration, one using housing starts data while the other covers U.S. post WWII individual production.

Nous présentons une classe générale de modèles non-linéaires avec changement de régime Markovienne. Les modèles proposés permettent d'avoir une structure périodique pour la chaîne de Markov ainsi que des effets saisonniers dans chaqu'un des régimes. La classe de structure proposée permet d'avoir des interdépendences entre les fluctuations saisonnières, les cycles d'affaire et la composante de croissance. Une méthode Baysienne basée sur le principe de l'échantillonage de Gibbs est utilisée pour estimation et interférence. Deux exemples empiriques sont fournis, un premier utilisant des séries de mise en chantier de maisons, tandis que le second couvre la production industrielle aux États-Unis.

Keywords: Markov Switching, Periodic Models, Seasonality, Gibbs Sampler
JEL: C11, C15, C22
Mots clés : modèles à changement de régime, structure périodique, saisonnalité, échantillonage de Gibbs

[†] C.R.D.E., Université de Montréal, and CIRANO.
[‡] Graduate School of Business, The University of Chicago.

1. Introduction

Modeling seasonality in nonlinear time series analysis is a relatively unexplored area. In this paper, we present a class of nonlinear time series models for seasonal data. Seasonal phenomena considered are not limited to linear characteristics such as deterministic mean shifts or peaks in the spectral decomposition at the seasonal frequency and its harmonics. The time series models considered can, for instance, predict that, say booms in housing starts are less likely to take off in the winter, that stock market crashes, economic recoveries, etc. appear less likely to occur during certain times of the year. It may also produce asymmetries in seasonal patterns and other nontrivial dependencies between seasonal, cyclical and long-term patterns in the data. Our analysis builds on a class of models in nonlinear time series analysis gaining considerable interest in recent years. It consists of a stochastic regime-switching structure driven by a hidden Markov process with a finite number of regimes. In econometrics, for instance Quandt (1960) and Goldfeld and Quandt (1973) proposed switching regression models, while Neftci (1984) and particularly Hamilton (1989) further developed such models as tools to investigate asymmetries in the cyclical behavior of macroeconomic aggregate series. More recently, McCulloch and Tsay (1993, 1994) introduced a more general class of Markov switching models which allows for state dependent AR polynomials, random variance shifts, and transition probabilities depending on a set of exogenous variables. Ghysels (1992, 1994) in work related to Hamilton (1989), proposed the use of periodic Markov switching structures to model the aforementioned nontrivial dependencies between different types of cycles such as seasonal and business cycles. A periodic Markov regime-switching structure is one where the Markov chain is nonhomogeneous, with the time variation of the transition probabilities being purely periodic, i.e., having the same transition scheme each year during a particular quarter or month, etc. In this paper, we exploit results of McCulloch and Tsay (1993, 1994) and Ghysels (1992, 1994) to propose a general Markov switching structure appropriate for seasonal time series.

The traditional maximum likelihood approach to Hamilton's original model, where a two-state Markov chain governs an intercept shift and the intertemporal dynamics are shaped by a fixed AR polynomial, is computationally demanding. In part, to overcome the computational burden, both Albert and Chib (1991) and McCulloch and Tsay (1994) independently proposed the more flexible environment of the Gibbs sampler, adopting a Bayesian approach to estimation and inference to overcome some of the computational difficulties.

In this paper, we also adopt a Bayesian approach to estimation using the Gibbs sampler as a simulation tool. This approach is particularly suited as a classical estimation of periodic Markov chain models often result in boundary parameter estimates, see Ghysels (1992, 1994) for further discussions. It is, in fact, quite useful to exploit extra-sample information regarding switching probabilities within a Bayesian framework. In this paper, we provide a detailed discussion of this and other issues. Section 2 provides a detailed description of our model and estimation procedure. Section 3 presents examples and Section 4 concludes.

2. A General Periodic Regime-Switching Model

Consider a discrete-time Markov chain process $\{\xi_t\}$ with two states, namely $\xi_t \in \{1, 2\}$. The process is a periodic Markov regime-switching process if its realizations ξ_t are governed by a periodic probability scheme with periodicity s and transition matrices $P(v)$ defined by

$$P(v) = \begin{bmatrix} 1-\epsilon_{v1} & \epsilon_{v1} \\ \epsilon_{v2} & 1-\epsilon_{v2} \end{bmatrix}, \quad v = 1, \cdots, s. \tag{2.1}$$

In other words, ξ_t is a two-state periodic Markov process with transition matrices $P(v)$, where v denotes the season at time index t. In economic applications, s may denote the number of seasons in a year and the states may represent the status of an economy. In this paper, we shall refer states as regimes of the model.

A time series $\{y_t\}$ follows a periodic regime-switching model if its evolution is governed by a hidden periodic Markov process ξ_t. Specifically, y_t satisfies the model

$$y_t = \begin{cases} X_t'\beta_{v,1} + Y_t'\phi_{v,1} + a_{v,1,t} & \text{if } \xi_t = 1 \\ X_t'\beta_{v,2} + Y_t'\phi_{v,2} + a_{v,2,t} & \text{if } \xi_t = 2 \end{cases} \tag{2.2}$$

where $X_t = (x_{1t}, \cdots, x_{kt})'$ is a set of exogenous regressors, including possibly a constant and some indicator variables, $Y_t = (y_{t-1}, \cdots, y_{t-p})'$ is a set of lagged dependent variables, and $\{a_{v,i,t}\}$ is a sequence of independent Gaussian random variates with mean zero and variance $\sigma^2_{v,i}$. In (2.2), the innovations $\{a_{v,i,t}\}$ are independent for different v and i. This model is a generalization of the Markov switching model of McCulloch and Tsay (1994) by introducing season-dependent transition matrices in (2.1). It is also related to the models considered by Tyssedal and Tjϕstheim (1988) and by Hamilton (1989).

If X_t contains a constant and y_t is stationary, then the stochastic structure of (2.2) allows for random mean shifts which vary according to regime and season, producing asymmetries in seasonal mean shifts. This property is similar to that of using seasonal dummy effects that differ across the business cycle. If X_t contains an indicator variable for a particular season and the associated parameter does not depend on v, then the model can be used to estimate seasonal effects which may be common to all regimes or depend on the regimes. Moreover, the AR polynomial $\phi_{v,i}$ in (2.2) may depend on the season, hence producing periodic time-variation in the dynamic of the process as well as in the regime transition. In general, the structure of model (2.2) is fairly rich. It enables the regime swiching to occur with higher probability in certain times of the year. The mean-shifts may depend on the regime and the season, and so may the dynamic structure of the system. Strictly speaking, however, the model in (2.2) is not identifiable, because the labels of the regimes and the parameter values of the submodels are interchangeable. Consequently, some constraints are needed to render the model identifiable. For further discussion, we follow the approach of McCulloch and Tsay (1994) by classifying the parameters of model (2.2) into three categories. The parameters in the first category are regime-invariant, that is, these parameters are the same for both regimes. The second category consists of parameters that are unique to regime i and season v, where $i = 1$, 2 and v denotes the season of time t. Finally, the third category contains constrained parameters for regime i and season v. The constraints specify prior information about how the parameters differ between the two regimes. These constrained parameters are used in applications to identify the proposed periodic regime-switching model.

Specifically, we partition each of the two regressor vectors of (2.2) into three subvectors, namely X_t by $(X'_{ct}, X'_{rt}, X'_{gt})'$ and Y_t by $(Y'_{ct}, Y'_{rt}, Y'_{gt})'$, and write the model as

$$y_t = \begin{cases} X'_{ct}\beta^c_v + Y'_{ct}\phi^c_v + X'_{rt}\beta^r_{v,1} + Y'_{rt}\phi^r_{v,1} + X'_{gt}\beta^g_{v,1} + Y'_{gt}\phi^g_{v,1} + a_{v,1,t} & \text{if } \xi_t = 1 \\ X'_{ct}\beta^c_v + Y'_{ct}\phi^c_v + X'_{rt}\beta^r_{v,2} + Y'_{rt}\phi^r_{v,2} + X'_{gt}\beta^g_{v,2} + Y'_{gt}\phi^g_{v,2} + a_{v,2,t} & \text{if } \xi_t = 2 \end{cases} \tag{2.3}$$

with v denoting the season of time t, where the parameter vectors are partitioned accordingly such that β^c_v and ϕ^c_v are the same for both regimes, but may depend on the season v, $\beta^r_{v,i}$ and $\phi^r_{v,i}$ are unique to regime i and season v, and $\beta^g_{v,i}$ and $\phi^g_{v,i}$ are constrained parameters for regime i and season v. The constraints are typically inequality constriants to separate the regimes and/or seasons. For example, if y_t is the growth rate of the U.S. quarterly real GNP and the regimes represent "expansions" and "contractions" of U.S. economy, then one might set an inequality constraint on the parameter of the constant term in X_t so that regime-1 has higher mean level. By so doing, one identifies regime-1 as expansion periods which have higher average growth rate. For further details, see McCulloch and Tsay (1994). If desirable, the innovational variances $\sigma^2_{v,i}$ can also be constrained.

The model in (2.2) can be partitioned in other ways than that in (2.3).

However, for the purpose of this paper, it suffices to consider model (2.3). In what follows, we consider a Bayesian analysis of such a two-state periodic regime-switching model by using the Gibbs sampler.

2.1. A Bayesian analysis

For simplicity, we assume that the first p observations $y_1, \cdots, y_p$ are given, where p is the maximum past lagged variable in Y_t. Let the observational vector be $y = (y_1, \cdots, y_n)'$ and the state vector $\xi = (\xi_{p+1}, \cdots, \xi_n)'$, where n is the sample size. Group the parameters in (2.3) as

$$\Psi^c = [(\Psi_1^c)', \cdots, (\Psi_s^c)']', \ \Psi^g = [(\Psi^g_{1,1})', (\Psi^g_{1,2})', \cdots, (\Psi^g_{s,1})', (\Psi^g_{s,2})']', \ \Psi_i^r = [(\Psi^r_{1i})', \cdots, (\Psi^r_{si})']' \tag{2.4}$$

where $\Psi_v^c = [(\beta_v^c)', (\phi_v^c)']'$, $\Psi^g_{v,i} = [(\beta^g_{v,i})', (\phi^g_{v,i})']'$, and $\Psi^r_{vi} = [(\beta^r_{v,i})', (\phi^r_{v,i})']'$ for $i = 1, 2$ and $v = 1, \cdots, s$. Thus, Ψ^c is the collection of parameters common to both regimes, Ψ^g denotes constrained parameters, and Ψ_i^r contains all un-constrained parameters that are unique to regime i.

Besides the coefficient parameters in (2.4), model (2.3) also invoves the innovational variances $\sigma^2_{v,i}$, the transition probabilities ϵ_{vi}, and the state configuration ξ. Any conventional statistical analysis of such a model would require a substantial amount of computing and in many cases become infeasible. To overcome this problem, we adopt a Bayesian approach via the Gibbs sampler. The Gibbs sampler is a recent development in the statistical literature and has been found to be useful in solving complicated statistical problems. See Casella and George (1992) for an introduction to Gibbs sampling and McCulloch and Tsay (1994) for its use in Markov switching models. Briefly speaking, the Gibbs sampler is a stochastic substitution procedure that enables us to make joint statistical inference from a set of conditional distributions of parameters given all the other parameters in the model.

Let H be the collection of all parameters in model (2.3) and $p(w|H - w, y)$ be the conditional posterior distribution of the parameter w given the data y and all the other parameters of the model, where $H - w$ denotes all the parameters of the model except w. The Gibbs sampler for the proposed model involves drawing random variates sequentially from the following conditional posterior distributions:

1. $p(\Psi^c|H - \Psi^c, y)$.
2. $p(\Psi_i^r|H - \Psi_i^r, y)$ for $i = 1, 2$.
3. $p(\Psi^g_{v,i}|H - \Psi^g_{v,i}, y)$ for $i = 1, 2$ and $v = 1, \cdots, s$.
4. $p(\xi_t|H - \xi_t, y)$ for $t = p + 1, \cdots, n$.

5. $p(\sigma^2_{v,i}|H - \sigma^2_{v,i}, y)$ for $i = 1, 2$ and $v = 1, \cdots, s$.

6. $p(\epsilon_{vi}|H - \epsilon_{vi}, y)$ for $i = 1, 2$ and $v = 1, \cdots, s$.

More specifically, the proposed Bayesian analysis of model (2.3) via the Gibbs sampler consists of the following steps:

1. Specify some prior distributions and some initial values for all the parameters in the model.

2. Draw random variates sequentially according to the conditional posterior distributions listed above. Once a realization of a parameter is drawn, it is treated immediately as the value of that parameter in the subsequent drawings of other parameters. The collection of all realizations in a pass through the conditional posterior distributions listed above is called a Gibbs iteration.

3. Iterate the Gibbs sampler for $M + N$ times. Discard the first M iterations, but keep the realizations of the last N iterations to form a Gibbs sample of size N on which statistical inference of the model can be made.

Under some mild regularity conditions, e.g. Geman and Geman (1984) and Tierney (1993), the Gibbs iterations form a Markov chain and, by ergodic theory, the sample joint distribution of the Gibbs sample converges weakly to the joint distribution of all the parameters. Therefore, marginal distributions of parameters of interest can easily be deduced from the Gibbs sample for making inference. The key condition for convergence is that the Markov chain of the Gibbs iteration is irreducible, which is true for the Markov switching models considered in the paper provided that the model is identifiable. In practice, different initial parameter values and different numbers of iterations should be used to check the convergence. Similarly, different prior specifications should be used to study the prior sensitivity in making inference.

It remains to complete the conditional posterior distributions listed above. To this end, we use proper conjugate priors, namely

$$\Psi^c \sim N(\Psi^c_0, A_c^{-1}),\ \Psi^r_i \sim N(\Psi^r_{i,0}, A_{r,i}^{-1}),\ \sigma^2_{vi} \sim \frac{u(v,i)\lambda_{vi}}{\chi^2_{u(v,i)}},\ \epsilon_{vi} \sim \mathrm{Beta}(\gamma_{v,i,1}, \gamma_{v,i,2}) \tag{2.5}$$

where $N(\mu, \Sigma)$ denotes a multivariate Gaussian distribution with mean μ and covariance matrix Σ, χ^2_u denotes chi-square dsitribution with u degrees of freedom, and $\mathrm{Beta}(\gamma_1, \gamma_2)$ is a beta-distribution with parameters γ_1 and

γ_2. For the constrained parameter Ψ^g, we employ componentwise inequality constraints,

$$\beta^g_{v1,j} < (\text{or } >)\beta^g_{v2,j} + \eta_{x,j} \quad \text{and} \quad \phi^g_{v1,j} < (\text{or } >)\phi^g_{v2,j} + \eta_{y,j}$$

where $\eta_{x,j}$ and $\eta_{y,j}$ are given constants and $\beta^g_{vi,j}$ and $\phi^g_{vi,j}$ denote the j-th element of β^g_{vi} and ϕ^g_{vi}, respectively. Let η be the collection of all constraint constants $\eta_{x,j}$ and $\eta_{y,j}$. The prior distribution of Ψ^g is then

$$\Psi^g \sim N(\Psi^g_0, A_g^{-1})I(\eta) \tag{2.6}$$

where $I(\eta)$ is an indicator function such that $I(\eta) = 1$ if Ψ^g satisfies the inequality constraints given by η. In the prior specifications in (2.5) and (2.6), all the hyper-parameters Ψ^c_0, $\Psi^r_{i,0}$, Ψ^g_0, A_c, $A_{r,i}$, A_g, $u(v,i)$, λ_{vi}, $\gamma_{v,i,j}$ and η are assumed to be known. If necessary, one can treat these hyper-parameters as parameters governed by yet another level of prior distributions. As mentioned earlier, sensitivity analysis of these hyper-parameters is an integral part of the proposed Bayesian analysis. Finally, we assume that all the prior distributions in (2.5) and (2.6) are independent of each other.

With the conjugate priors in (2.5) and (2.6), the necessary conditional posterior distributions can be obtained by traditional Bayesian techniques, e.g. DeGroot (1970). Details of those conditional posterior distributions can be found in McCulloch and Tsay (1994) with some modifications.

Some remarks are in order in our implementation of the Gibbs sampling. First, instead of individual state ξ_t, we draw δ states jointly, say $\xi_{t,\delta} = (\xi_t, \xi_{t+1}, \cdots, \xi_{t+\delta-1})$. This modification serves two purposes: (a) it can speedup the convergence of the Gibbs sampler, because realizations of adjacent states are often dependent, and (b) by varying δ, we can check the convergence of the sampler. Second, the constrained parameters in Ψ^g are drawn component by component. This enables us to check the constraints easily. Third, the drawing of innovational variances σ^2_{vi} can be simplified if the process is homogenous across the rgimes and/or seasons.

3. Examples

3.1. Monthly Housing Starts

The first series we consider consists of monthly housing starts for single family homes from 1964 to 1991. Figure 1 is a time series plot of the data which clearly were seasonally unadjusted. The series is a closely watched leading indicator, as increased activities on housing starts are often a prelude to economic recovery.

To implement the modelling procedure described in section 2, we first choose the seasonal structure for the switching probabilities (ϵ_{vi}). We use two seasons, with season 1 being the months {January, February, November, December}, and season 2 encompassing the rest of the months. With this specification we can investigate whether the switching mechanism is different in the winter months, a plausible hypothesis for the housing industry.

For our switching model we let the current value y_t depend on an intercept, a dummy variable set to 1 for the winter months as defined by our choice of seasons, and lagged values with lags of 1, 2, and 12. The dummy and lag 12 are included to capture the seasonal structure of the data. We must also specify, for each explanatory variable, which of the three classes discussed in section 2 its coefficient belongs to. To identify the two states, we let the intercept be regime dependent with the constraint that the intercept in regime 1 be larger than the intercept in regime 2 by 0.1. The rest of the coefficients and the residual variance are constrained to have the same value in both regimes. In the notation of equation 2.3, we have X_{gt} consists of a vector of ones, X_{ct} consists of the seasonal dummy variable, and Y_{ct} consists of the values of y_t lagged 1, 2, and 12 periods. Intuitively we now think of seasons 1 and 2 as "winter" and "summer" and regimes 1 and 2 as "high" and "low" housing activities.

The prior for the two intercepts (one for each regime) is the bivariate normal with 0 mean and covariance matrix equal to four times the identity, conditioned on the restriction given above. All other coefficients have independent $N(0,1)$ priors. The residual variance has prior $\sigma^2 \sim \frac{\nu\lambda}{\chi^2_\nu}$ with $\nu = 5$ and $\lambda = 0.2$. Finally, there are the four ϵ_{vi} where v indexes the two seasons and i indexes the two regimes. Recall that ϵ_{vi} denotes the probability of leaving regime i in season v. The four ϵ_{vi} are independent Beta(2,15). The first boxplot in figure 3 depicts this common prior. All of the priors have been chosen in an attempt to make them spread out, without supporting improbable values. For example, making the autoregressive coefficients standard normal would seem to cover the range of likely values. We do not expect the coefficient for y_{t-1} to be 3! Choosing the prior for the dummy coefficient and residual variance is less obvious and our choices reflect a fair amount of playing around with the model.

In order to compute the joint posterior distribution, we ran the Gibbs sampler for $M = 500$ initial iterations and kept the results for a subsequent $N = 10,000$ iterations. We compared the output obtained from different runs and found the results to be very stable.

Figure 2 plots $P(\xi_t = 2 \mid Y, X)$, the posterior probability that time t is in regime 2 versus time index t. Note that although our data start in 1964 the plot starts in 1965 because of the lagged variables used. The plot shows for example, that there is strong evidence that housing starts was in regime 2 in the mid 70's and in the early 80's. Except for the second half of the 80's the results are quite strong in that the probabilities are close to zero or one.

Is the propensity to switch regimes related to the season? Figure 3 uses boxplots to display the prior and posterior distributions of the ϵ_{vi}. The first boxplot is based on 1000 draws from the common Beta(2,15) prior. The remaining four boxplots are based on our 10000 draws from the posterior distribution of the ϵ_{vi}. The boxplots suggest that there is strong evidence that the switching probabilities are related to the seasons. In the first season (winter) it appears that ϵ_{11}, the probability of leaving regime 1, is substantially smaller than ϵ_{12}, the probability of leaving regime 2. In the second season (summer) the ordering is reversed with ϵ_{22} concentrated on the smaller values. However, some care in needed in looking at the boxplots. In particular, while the posterior for ϵ_{12} seems to support value much larger than those of ϵ_{11}, we note that the posterior of ϵ_{12} is very similar to the prior so that what may really be going on here is that there is little information in the sample about ϵ_{12}.

Table 3.1 shows the median, 5%, and 95% quantiles for the marginal posteriors of the model parameters. Note that the posterior probability that the difference between the intercepts in the two regimes is in the interval (.14, .19) is 90%. Although both intercepts vary over a similar range, they are highly dependent and the marginal distribution of their difference is concentrated near .17. The differences in the transition probabilities we saw in figure 3 are also evident in the table. For example, the posterior median of $\epsilon_{1,2}$ is .12 while the posterior median of $\epsilon_{2,2}$ is .031.

3.2. Quarterly Industrial Production

Our second series is the index of U.S. quarterly Industrial Production from the first quarter of 1947 to the final quarter of 1991. The data were seasonally adjusted and obtained from the Citibase data system. Our initial step in the analysis is to difference the data. The differenced series is displayed in figure 4.

To investigate the possibility of seasonal switching probabilities we specify two seasons with the first season being the first and fourth quarters and the second season being the second and third quarters. Again, we can roughly think of our two seasons as being "winter" and "summer".

The switching model includes an intercept and lagged production indicies at lags 1, 2, 5, and 8. We used the intercept to define the regimes by letting it be regime dependent with the contraint that the intercept in regime 1 be larger than the intercept in regime 2 by at least 0.5. The four

Table 3.1: Housing Starts: Selected Quantiles of Marginal Posteriors of Parameters

parameter	median	.05 quantile	.95 quantile
intercept, regime 1	3.94	3.57	4.44
intercept, regime 2	3.77	3.41	4.27
difference in intercepts	.17	.14	.19
ar1 coefficient	.69	.60	.78
ar2 coefficient	-.18	-.25	-.10
seasonal dummy	-.21	-.24	-.18
ar12 coefficient	.15	.09	.20
σ	.12	.11	.13
$\epsilon_{1,1}$	.031	.006	.086
$\epsilon_{1,2}$	.120	.034	.245
$\epsilon_{2,1}$	.083	.036	.165
$\epsilon_{2,2}$	.031	.007	.081

coefficients of the lagged y_t values were allowed to be state dependent without any constraint. In the notation of equation 2.3, we have X_{gt} consists of a vector of ones, and Y_{rt} consists of y_{t-i} for $i = 1$, 2, 5, and 8.

The prior used for the intercept and autoregressive coefficients is the same as in the previous example (standard normal) with corresponding coefficients in the two regimes being independent. The residual variance is constrained to be the same in the two regimes with $\nu = 5$ and $\lambda = 1$. The ϵ_{vi} are independent Beta(5,15) as depicted in the first boxplot in figure 6.

Figure 5 plots $P(\xi_t = 2 \mid Y, X)$, the posterior probability that time t is in regime 2 versus t. The majority of quarters are in the first regime with high probability. While this plot is markedly different from the corresponding plot in the previous example, note that they are somewhat in accordance in that high probabilities for state 2 occur around 1970, the mid seventies, early eighties, and early nineties.

Figure 6 (analogous to figure 3) depicts the marginal posteriors of the ϵ_{vi}. This time the switching pattern is similar in both seasons. In both seasons, the data has moved the distribution of ϵ_{v1} towards smaller values and moved the distribution of ϵ_{v2} towards larger values relative to the common prior distribution of the ϵ_{vi}. Thus, as clearly suggested in figure 5, it is generally much easier to leave regime 2 than regime 1. There is however, the suggestion that regime 1 is "stickier" in season 1 (winter) than in season 2. The median value of the 10,000 draws of ϵ_{11} is .08 while it is .14 for ϵ_{21}.

Table 3.2: Differenced Industrial Production: Selected Quantiles of Marginal Posteriors of Parameters

parameter	median	.05 quantile	.95 quantile
intercept, state 1	.81	.64	.97
intercept, state 2	-1.42	-1.94	-.84
ar1 coefficient, regime 1	.35	.24	.47
ar1 coefficient, regime 2	.75	.37	1.08
ar2 coefficient, regime 1	-.15	-.25	-.04
ar2 coefficient, regime 2	-.28	-.74	.22
ar5 coefficient, regime 1	-.17	-.26	-.08
ar5 coefficient, regime 2	.83	.31	1.34
ar8 coefficient, regime 1	-.16	-.26	-.06
ar8 coefficient, regime 2	-.60	-.94	-.29
σ	.79	.71	.88
$\epsilon_{1,1}$	.09	.038	.157
$\epsilon_{1,2}$	.342	.199	.499
$\epsilon_{2,1}$	.140	.082	.222
$\epsilon_{2,2}$	.313	.168	.491

Table 3.2 shows the median, 5%, and 95% quantiles for the marginal posteriors of the model parameters. In this example the intercepts are markedly different in the regimes. Also there seems to be strong evidence that the dynamic structure of the model in the two regimes is different (note for example the lag 5 AR coefficients).

4. Concluding Remarks

In this paper we have proposed a general class of periodic Markov regime-switching models and used Gibbs sampler to analyze such models. Two real economic examples are used to illustrate the application of the models. We found evidence that the asymmetric transition probabilities between regimes are also season dependent both for the U.S. monthly housing starts and quarterly industrial production index. This suggests that for both data sets considered their behavior in the winter is different from that in the summer. In particular, the magnitudes of transition probabilities reverse between winter and summer for the monthly housing starts series.

References

Albert, J. and Chib, S. (1993), "Bayesian Inference of Autoregressive Time Series With Mean and Variance Subject to Markov Jumps," *Journal of Business & Economic Statistics*, 11, 1-15.

Casella, G., and George, E.I. (1992), "Explaining the Gibbs Sampler," *The American Statistician*, 46, 167-174.

DeGroot, M. (1970), *Optimal Statistica Decisions*, McGraw-Hill, New York.

Gelfand, A. E. and Smith, A. F. M. (1990), "Sampling-Based Approaches to Calculating Marginal Densities," *Journal of the American Statistical Association*, 85, 398-409.

Geman, S. and Geman, D. (1984), "Stochastic Relaxation, Gibbs Distributions and the Bayesian Restoration of Images," *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6, 721-741.

Ghysels, E. (1992), "A Time Series Model with Periodic Stochastic Regime Switching," Discussion Paper, C.R.D.E., Université de Montréal.

Ghysels, E. (1994), "On the Periodic Structure of the Business Cycles," *Journal of Business and Economic Statistics*, 12, 289-299.

Goldfeld, S. M. and Quandt, R.E. (1973), "A Markov Model for Switching Regressions," *Journal of Econometrics*, 1, 3-16.

Hamilton, J. D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357-384.

McCulloch, R. E. and Tsay, R. S. (1994), "Statistical Inference of Macroeconomic Time Series via Markov Switching Models," *Journal of Time Series Analysis*, forthcoming.

Neftci, S N. (1984), "Are Economic Time Series Asymmetric Over the Business Cycle?," *Journal of Political Economy*, 92, 307-328.

Quandt, R. E. (1960), "Tests of the Hypothesis that a Linear Regression System Obeys Two Separate Regimes," *Journal of the American Statistical Association*, 55, 324-330.

Tierney, L. (1994), "Markov Chains for Exploring Posterior Distributions," *Annals of Statistics*, forthcoming.

Tyssedal, J. S. and Tjøstheim, D. (1989), "An Autoregressive Model with Suddenly Changing Parameters and an Application to the Stock Market," *Applied Statistics*, 37, 353-369.