Homework 2

Due June 1st at 1pm.

For $t = 1, \ldots, n$, let us consider the following version of a local-level dynamic linear model:

$$
\begin{aligned}
y_t &= x_t + v_t \\
x_t &= \phi x_{t-1} + w_t \qquad\qquad w_t \overset{iid}{\sim} N(0, \tau^2),
\end{aligned}
$$

with $x_0 \sim N(m_0, C_0)$.

We consider two possible structures for $v_1, \ldots, v_n$:

$$
\begin{aligned}
\mathcal{M}_1 &: \ v_t \sim \pi N(0, \sigma^2) + (1 - \pi) N(0, \kappa^2 \sigma^2), \\
\mathcal{M}_2 &: \ v_t \sim t_\nu(0, \sigma^2),
\end{aligned}
$$

where $\pi \in (0, 1)$, $\kappa > 0$ and $\nu > 0$.

Data augmentation

Notice that $v_t \sim \pi N(0, \sigma^2) + (1-\pi) N(0, \kappa^2 \sigma^2)$ can be rewritten as

$$v_t | \lambda_t \sim N(0, \sigma^2_{\lambda_t}) \quad \text{and} \quad \lambda_t \overset{iid}{\sim} Ber(\pi)$$

where $\sigma_t^2 = \sigma^2_{\lambda_t}$ with $\sigma_0^2 = \sigma^2$ and $\sigma_1^2 = \kappa^2 \sigma^2$.

Similarly, $v_t \sim t_\nu(0, \sigma^2)$ can be rewritten as

$$v_t | \lambda_t \sim N(0, \lambda_t \sigma^2) \quad \text{and} \quad \lambda_t \sim IG(\nu/2, \nu/2).$$

Conditionally on $\{\lambda_t\}_{t=1}^n$, both models are standard normal dynamic linear models (NDLMs).

Prior

We will assume that

$$\begin{aligned} p(\sigma^2, \tau^2, \pi, \kappa | \mathcal{M}_1) &= p(\sigma^2)p(\tau^2)p(\pi)p(\kappa) \\ p(\sigma^2, \tau^2, \nu | \mathcal{M}_2) &= p(\sigma^2)p(\tau^2)p(\nu) \end{aligned}$$

where

$$\sigma^2 \sim IG(\nu_0/2, \nu_0\sigma_0^2/2)$$

$$\tau^2 \sim IG(\eta_0/2, \eta_0\tau_0^2/2)$$

$$\pi \sim U(0,1)$$

$$\kappa \sim IG(a,b)$$

$$\nu \sim \text{uniform on } \{1,2,\ldots,m\}$$

MCMC

Let $y^n = (y_1, \ldots, y_n)$ and $x^n = (x_1, \ldots, x_n)$.
Let $\theta_1 = (x_0, \phi, \sigma^2, \tau^2, \pi, \kappa)$ and $\theta_2 = (x_0, \phi, \sigma^2, \tau^2, \nu)$.

Derive MCMC schemes to sample from

$$p(x^n, \theta_1 | y^n, \mathcal{M}_1) \quad \text{and} \quad p(x^n, \theta_2 | y^n, \mathcal{M}_2),$$

by taking into account that the latent variables $\{\lambda_t\}_{t=1}^n$ facilitate the derivation of easy-to-sample full conditional distributions.

Simulation exercises

In order to test both algorithms, simulate two sets of $n = 200$ observations, one from $\mathcal{M}_1$ and one from $\mathcal{M}_2$, with the following specifications:

$$\mathcal{M}_1 : \theta_1 = (0.0, 0.9, 1.0, 1.0, 0.9, 2.0)$$

$$\mathcal{M}_2 : \theta_2 = (0.0, 0.9, 1.0, 1.0, 4.0)$$

Run your MCMC schemes for 20,000 draws and discard the first half as burn-in. Use the true values as initial values for the fixed parameters θ_1 and θ_2.

Suggestion

First run your code assuming the fixed parameters are known since

$$p(x^n|\theta, y^n, \mathcal{M}_i), \qquad i = 1, 2,$$

are the more involving full conditionals. Well, not really! You can use the augmented latent variables $(\lambda_1, \ldots, \lambda_n)$ to neatly derive quite standard FFBS schemes.

Then breaking

$$p(\theta|x^n, y^n, \mathcal{M}_i), \qquad i = 1, 2,$$

into univariate full conditionals should be straightforward.

Presentations

Two of the three 3-person groups will be selected at the beginning of the class on June 1st to present the results of the two models (one model each group). The third group is off the hook, but that will be known only right before the presentations.

Be ready to describe your simulation and your code.

Each presentation will last at most 15 minutes.