Second homework assignment

PhD in Business Economics — Course: Econometrics III

Professor: Hedibert Freitas Lopes — Due date: 10h30, February 11th, 2020.

You can work individually or in pairs.

Use, preferably, `Rmarkdown` (via `RStudio`) to produce your report in PDF or HTML.

Problem: Fitting an ARMA(2,1) model

Let us assume that the data $\mathcal{D}_n = \{y_1, \ldots, y_n\}$ follows an ARMA(2,1) model

$$y_t = \phi_0 + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \varepsilon_t + \theta\varepsilon_{t-1}$$

where $\varepsilon_1, \ldots, \varepsilon_n$ are i.i.d. $N(0, \sigma^2)$. To simplify the problem, let us assume that $\varepsilon_0 = 0$ and that $y_0 = y_{-1} = 0$. Therefore, it is easy to see that $\{\varepsilon_1, \ldots, \varepsilon_n\}$ are deterministically obtained from $\{\varepsilon_0, y_0, y_{-1}, \phi_0, \phi_1, \phi_2, \theta, \sigma^2\}$ and the data $\mathcal{D}_n$:

$$\begin{aligned}
\varepsilon_1 &= y_1 - \phi_0 \\
\varepsilon_2 &= y_2 - \phi_0 - \phi_1 y_1 - \theta\varepsilon_1 \\
\varepsilon_t &= y_t - \phi_0 - \phi_1 y_{t-1} + \phi_2 y_{t-2} - \theta\varepsilon_{t-1}, \qquad t = 3, \ldots, n.
\end{aligned}$$

Likelihood and prior. To avoid overloading the notation, let us drop $\{\varepsilon_0, y_0, y_{-1}\}$ in what follows. The likelihood function is, therefore,

$$\mathcal{L}(\Theta|\mathcal{D}_n) = \left(\frac{1}{2\pi\sigma^2}\right)^{-n/2} \exp\left\{-\frac{\sum_{t=1}^n \varepsilon_t^2}{2\sigma^2}\right\}.$$

where $\Theta = (\phi_0, \phi_1, \phi_2, \theta, \sigma^2)$. Let us assume the following prior for Θ:

$$p(\Theta) = p(\phi_0)p(\phi_1)p(\phi_2)p(\theta)p(\sigma^2),$$

and

$$\begin{aligned}
\phi_0 &\sim N(0, 100) \quad \text{and} \quad (\phi_1, \phi_2) \sim U(\mathcal{A}) \\
\theta &\sim U(0, 1) \quad \text{and} \quad \sigma^2 \sim IG(0, 0),
\end{aligned}$$

where $\mathcal{A} = \{(\phi_1, \phi_2) \in \mathbb{R}^2, \text{ such that } |\phi_2| < 1, \phi_1 + \phi_2 < 1, \phi_2 - \phi_1\}$ is the AR(2) stationary region.

Data. Inspired by the Canadian Lynx data presented and analysed in class, let us simulate $n = 100$ observations from the following ARMA(2,1) model

$$y_t = 1.0 + 1.4y_{t-1} - 0.8y_{t-2} + \varepsilon_t + 0.8\varepsilon_{t-1}, \qquad \varepsilon_t \sim N(0,1),$$

so $\Theta = (1.0, 1.4, -0.8, 0.8, 1.0)$. Feel free to use my script to simulate some data:

```
set.seed(12345)
n = 100
sig = 1.0
eps = rnorm(n,0,sig)
y = rep(0,n)
y[1] = 1.0+eps[1]
y[2] = 1.0+1.4*y[1]+0.8*eps[1]+eps[2]
for (t in 3:n)
  y[t] = 1.0+1.4*y[t-1]-0.8*y[t-2]+0.8*eps[t-1]+eps[t]
ts.plot(y)
print(arima(y,order=c(2,0,1)))
```

Questions:

1. What are the ML estimates (MLE) of $\phi_0, \phi_1, \phi_2, \theta, \sigma^2$? Use the `R` function `nlm` to minimize $-\mathcal{L}(\Theta|\mathcal{D}_n)$. Are the results similar to the ones from the `R` function `arima(y,order=c(2,0,1))`?

2. Use sampling importance resampling (SIR) to sample from the posterior of Θ

$$p(\Theta|\mathcal{D}_n) \propto \mathcal{L}(\Theta|\mathcal{D}_n)p(\Theta).$$

   (a) Compute posterior means, medians and 95% credibility interval for all components of Θ. Are posterior means (and medians) similar to their MLE counterparts?

   (b) Plot the posterior density of $\phi_1 + \phi_2$, i.e. $p(\phi_1 + \phi_2|\mathcal{D}_n)$. What is the posterior probability that $\phi_1 + \phi_2 > 0.8$?

   (c) Compute $E(y_{n+h}|\mathcal{D}_n)$ and $V(y_{n+h}|\mathcal{D}_n)$ for $h = 1, \ldots, 20$.

3. Repeat 2(a)-(c), but replace the SIR algorithm by an MCMC algorithm. Notice that the full conditionals of ϕ_0 and σ^2 are the only ones of known form (Gaussian and Inverse Gamma, respectively).