Exercises for "Bayesian Computation with R"

Rainer Hirk

WS 2021/2022

1. This problem deals with the analysis of the daily exchange rates Bitcoin/US dollar in year 2015. We fit an AR(1) process to the r_t time series of log returns:

$$r_t = \alpha + \beta r_{t-1} + \epsilon_t,$$

where $\epsilon_t \overset{iid}{\sim} N(0, \sigma^2)$ and t is the day indicator.

   (a) Load the data using the **quantmod** package (symbol: BTC/USD). Compute and plot the r_t series.

   (b) Using independent conjugate priors for α, β and $\tau = \sigma^{-2}$ derive the Gibbs sampler. Determine the full conditional distribution of each of the parameters α, β and τ (i.e., conditional on the data and the remaining parameters).

   (c) Implement the Gibbs sampler in R.

   (d) Compare the results of your Gibbs sampler to those obtained using package **rjags** or package **rstan** for this data example.

   (e) Which characteristic of the time series of log returns is not captured by the above model but can be seen from the plot? Which model can be used to overcome this "problem"?

   For the parameters of the priors use the following values (parametrization in R):

   - Normal distribution: $N(0, 10^6)$
   - Gamma distribution: Gamma$(10^{-3}, 10^{-3})$

2. Consider the series of log returns y_t for Apple, Inc for the period 2004-2014. Suppose $y_t \sim N(\mu, \sigma^2)$. Using a conjugate prior for $\mu \sim N(m, s^2)$, the following hypotheses are to be tested

$$H_1 : \mu \leq 0$$
$$H_2 : \mu > 0$$

   (Use non-informative hyperparameters for m and s^2. Assume the true variance known and $\sigma^2 = 0.05$).

   (a) Download the series of adjusted prices using the **quantmod** package (symbol `AAPL`) and compute the log returns.

   (b) Perform the frequentist test in R.

   (c) Compute the Bayes factor.

3. Look at the daily log returns in 2. Determine the number of daily changes in excess of 1% occuring in each month, i.e., increases and decreases larger than 1%.

   We assume that the data follows a Poisson distribution with the rate parameter θ.

(a) What is the likelihood function?

(b) What is the conjugate prior? Derive the posterior distribution.

(c) What is the Jeffrey's prior? What is the posterior for the rate parameter if Jeffrey's prior is used?

(d) Consider only the data for 2014. Try out different prior distributions, e.g., different parameter settings for the conjugate prior and Jeffrey's prior. Plot the prior distributions as well as the posterior distributions for comparison.

(e) Repeat 3d for the whole time series from 2004–2014. Are there differences in the distributions as the number of observations increases? Does the influence of the prior change?