Modern Bayesian Statistics
Part III: high-dimensional modeling
Example 1: Time-varying variance modeling

Hedibert Freitas Lopes[1]
hedibert.org

13^a aMostra de Estatística
IME-USP, October 2018

[1]Professor of Statistics and Econometrics at Insper, São Paulo.

Example 1: Time-varying variance modeling

Modeling Petrobrás' log-returns

Time span: 12/29/2000 - 12/31/2013 ($n = 3268$ days)

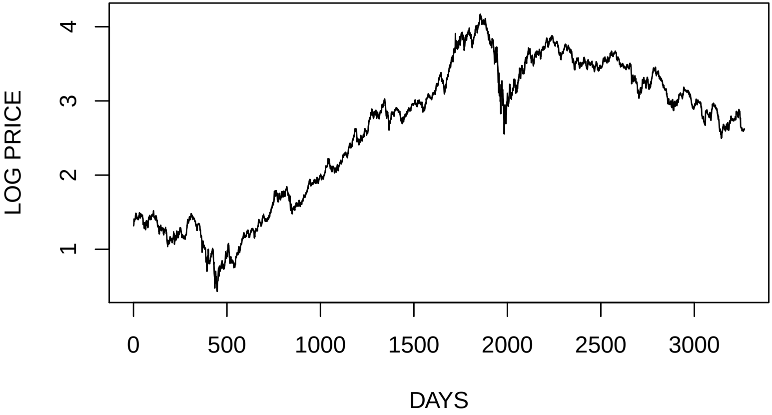

Scatterplot of y_{t-1} versus y_t

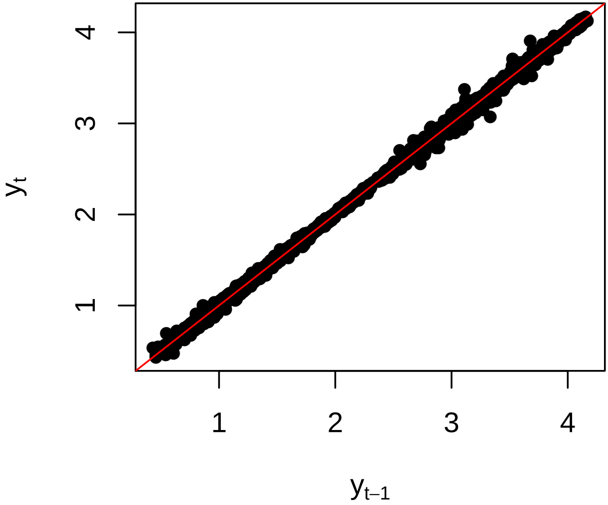

Log return: $r_t = y_t - y_{t-1} = \log(p_t/p_{t-1})$

Time span: 01/02/2001 - 12/31/2013 ($n = 3267$ days)

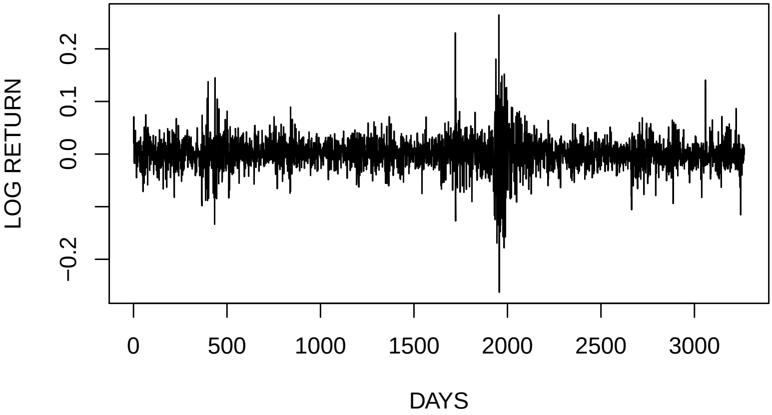

Histogram of r_t

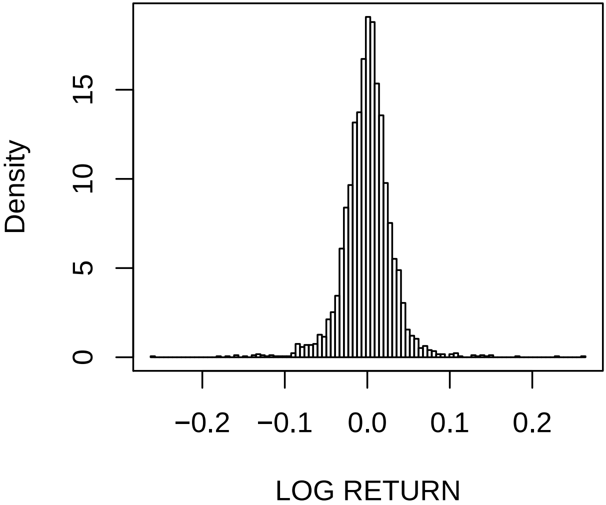

Training and testing samples

Years 2001-2006:
The first $n_0 = 1506$ days are used for prior specification.

Years 2007-2013:
The last $n = 1760$ days are used for posterior inference.

GARCH(1,1) with t errors

The GARCH(1,1) model with Student-t innovations:

$$
\begin{aligned}
r_t &\sim t_\nu(0, \rho h_t) \\
h_t &= \alpha_0 + \alpha_1 r_{t-1}^2 + \beta h_{t-1},
\end{aligned}
$$

where $\alpha_0 > 0$, $\alpha_1 \geq 0$ and $\beta > 0$.

We set the initial variance to $h_0 = 0$ for convenience.

We let $\rho = (\nu - 2)/\nu$ so that

$$
V(r_t | h_t) = \frac{\nu}{\nu - 2} \rho h_t = h_t.
$$

Prior

Let $\psi = (\alpha', \beta, \nu)'$ and $\alpha = (\alpha_0, \alpha_1)'$.

The prior distribution of ψ is such that

$$p(\alpha, \beta, \mu) = p(\alpha)p(\beta)p(\nu)$$

where

$$\begin{aligned} \alpha &\sim N_2(\mu_\alpha, \Sigma_\alpha) I_{(\alpha>0)} \\ \beta &\sim N(\mu_\beta, \Sigma_\beta) I_{(\beta>0)} \end{aligned}$$

and

$$p(\nu) = \lambda \exp\{-\lambda(\nu - \delta)\} I_{(\lambda>\delta)}$$

for $\lambda > 0$ and $\delta \geq 2$, such that $E(\nu) = \delta + 1/\lambda$.

Normal case: $\lambda = 100$ and $\delta = 500$.

bayesGARCH

`bayesGARCH`: Bayesian Estimation of the GARCH(1,1) Model with Student-t Innovations

```
bayesGARCH(r,mu.alpha = c(0,0),Sigma.alpha=1000*diag(1,2),
           mu.beta=0,Sigma.beta=1000,
           lambda=0.01,delta=2,control=list())
```

Paper: Ardia and Hoogerheide (2010) Bayesian Estimation of the GARCH(1,1) Model with Student-t Innovations. *The R Journal*, 2,41-47.

`http://cran.r-project.org/web/packages/bayesGARCH`

Example of `R script`

Recall that r_0 are Petrobras' returns for the first part of the data.

```
M0    = 10000       # to be discarded (burn-in)
M     = 10000       # kept for posterior inference
niter = M0+M

MCMC.initial = bayesGARCH(r0,mu.alpha=c(0,0),Sigma.alpha=1000*diag(1,2),
                          mu.beta=0,Sigma.beta=1000,lambda=0.01,delta=2,
                          control=list(n.chain=1,l.chain=niter,refresh=100))

draws = MCMC.initial$chain1

range = (M0+1):niter

par(mfrow=c(2,2))
ts.plot(draws[range,1],xlab="iterations",main=expression(alpha[0]),ylab="")
ts.plot(draws[range,2],xlab="iterations",main=expression(alpha[1]),ylab="")
ts.plot(draws[range,3],xlab="iterations",main=expression(beta),ylab="")
ts.plot(draws[range,4],xlab="iterations",main=expression(nu),ylab="")
```

MCMC output

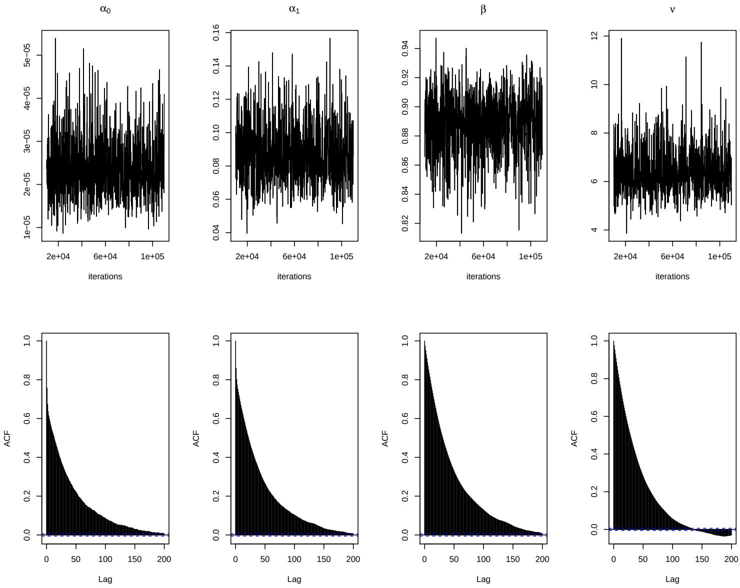

α0
α1
β
ν
5e−05
4e−05
3e−05
2e−05
1e−05
0.16
0.14
0.12
0.10
0.08
0.06
0.04
0.94
0.92
0.90
0.88
0.86
0.84
0.82
12
10
8
6
4
2e+04
6e+04
1e+05
iterations
ACF
1.0
0.8
0.6
0.4
0.2
0.0
0
50
100
150
200
Lag

Marginal posterior distributions

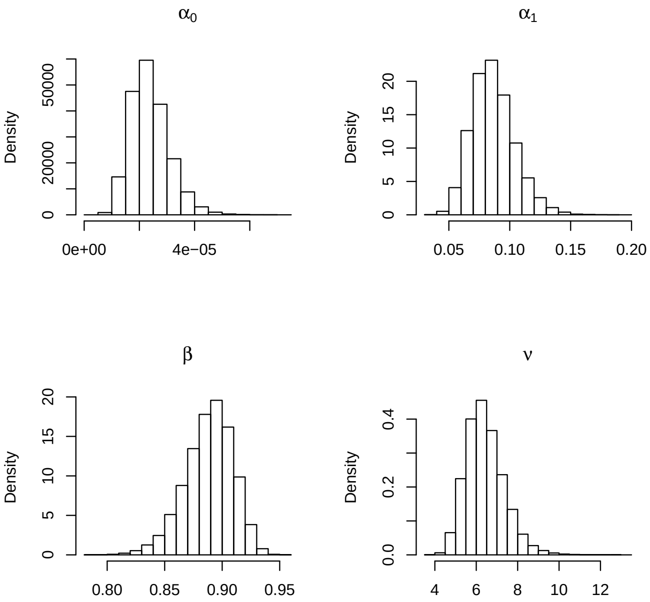

$p(\alpha_1 + \beta|\texttt{data})$

$$Pr(\alpha_1 + \beta > 1|\texttt{data}) = 0.0034$$

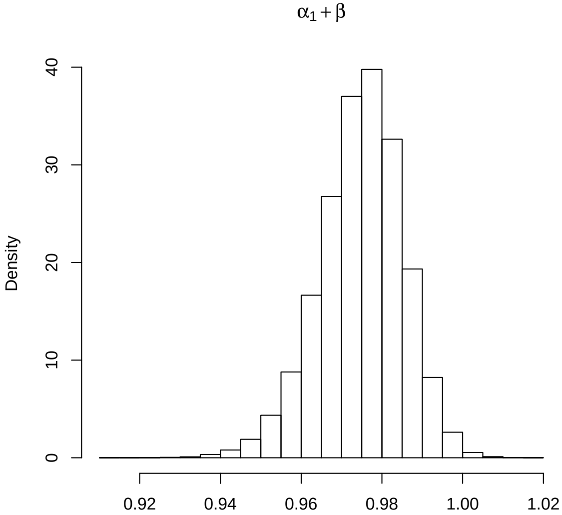

Quantiles from $p(h_t^{1/2}|\texttt{data})$

Percentiles 2.5%, 50% and 97.5% of $p(h_t^{1/2}|\texttt{data})$
Black vertical lines: r_t^2