



Munich Personal RePEc Archive

Bayesian Estimation of the GARCH(1,1) Model with Student-t Innovations in R

Ardia, David
University of Fribourg

20. September 2009

Online at <http://mpa.ub.uni-muenchen.de/17414/>
MPRA Paper No. 17414, posted 20. September 2009 / 21:16

Bayesian Estimation of the GARCH(1,1) Model with Student- t Innovations in R

David Ardia

Department of Quantitative Economics
University of Fribourg, Switzerland

Abstract

This paper presents the R package **bayesGARCH** which provides functions for the Bayesian estimation of the parsimonious but effective GARCH(1,1) model with Student- t innovations. The estimation procedure is fully automatic and thus avoids the time-consuming and difficult task of tuning a sampling algorithm. The usage of the package is shown in an empirical application to exchange rate log-returns.

Keywords: GARCH, Bayesian, MCMC, Student- t , R software.

1. Introduction

Research on changing volatility using time series models has been active since the pioneer paper by Engle (1982). From there, ARCH and GARCH type models grew rapidly into a rich family of empirical models for volatility forecasting during the 80's. These models are widespread and essential tools in financial econometrics and have, until recently, mainly been estimated using the classical Maximum Likelihood technique. The Bayesian approach offers an attractive alternative which enables small sample results, robust estimation, model discrimination and probabilistic statements on nonlinear functions of the model parameters.

The choice of the algorithm is the first issue when dealing with MCMC methods and it depends on the nature of the problem under study. The package **bayesGARCH** use the simulation procedure of Ardia (2008, Chapter 5), which relies on the M-H algorithm where some model parameters are updated by blocks. The densities are constructed from an auxiliary ARMA process for the squared observations. This methodology avoids the time-consuming and difficult task, especially for non-experts, of choosing and tuning a sampling algorithm.

The outline of the paper is as follows: The model specification and MCMC scheme are presented in Section 2. An empirical application is proposed in Section 3. Section 4 concludes.

2. Model, priors and MCMC scheme

A GARCH(1,1) model with Student- t innovations may be written via data augmentation (Ardia 2008, Chapter 5) as follows:

$$\begin{aligned} y_t &= \varepsilon_t \left(\frac{\nu-2}{\nu} \varpi_t h_t \right)^{1/2} \quad t = 1, \dots, T \\ \varepsilon_t &\stackrel{iid}{\sim} \mathcal{N}(0, 1) \quad \varpi_t \stackrel{iid}{\sim} \mathcal{IG} \left(\frac{\nu}{2}, \frac{\nu}{2} \right) \\ h_t &\doteq \alpha_0 + \alpha_1 y_{t-1}^2 + \beta h_{t-1}, \end{aligned} \quad (1)$$

where $\alpha_0 > 0$, $\alpha_1, \beta \geq 0$ and $\nu > 2$; y_t is a scalar dependent variable; $\mathcal{N}(\cdot, \cdot)$ denotes the standard Normal distribution; $\mathcal{IG}(\cdot, \cdot)$ denotes the Inverted Gamma distribution.

In order to write the likelihood function, we define the vectors $\mathbf{y} \doteq (y_1 \cdots y_T)'$ and $\boldsymbol{\alpha} \doteq (\alpha_0 \ \alpha_1)'$. We regroup the model parameters into $\boldsymbol{\psi} \doteq (\boldsymbol{\alpha}, \beta, \nu)$, the latent variables into the $T \times 1$ vector $\boldsymbol{\varpi} \doteq (\varpi_1 \cdots \varpi_T)'$, and define the augmented set of parameters $\Theta \doteq (\boldsymbol{\psi}, \boldsymbol{\varpi})$. Then, upon defining the $T \times T$ diagonal matrix:

$$\boldsymbol{\Sigma} \doteq \boldsymbol{\Sigma}(\Theta) = \text{diag}(\{\varpi_t \frac{\nu-2}{\nu} h_t(\boldsymbol{\alpha}, \beta)\}_{t=1}^T),$$

where $h_t(\boldsymbol{\alpha}, \beta) \doteq \alpha_0 + \alpha_1 y_{t-1}^2 + \beta h_{t-1}$, we can express the likelihood of parameter Θ as follows:

$$p(\mathbf{y} | \Theta) \propto (\det \boldsymbol{\Sigma})^{-1/2} \exp \left[-\frac{1}{2} \mathbf{y}' \boldsymbol{\Sigma}^{-1} \mathbf{y} \right]. \quad (2)$$

We propose the following proper priors on the parameters $\boldsymbol{\alpha}, \beta$ of the preceding model:

$$\begin{aligned} p(\boldsymbol{\alpha}) &\propto \phi_{\mathcal{N}_2}(\boldsymbol{\alpha} | \boldsymbol{\mu}_\alpha, \boldsymbol{\Sigma}_\alpha) \mathbb{I}\{\boldsymbol{\alpha} > \mathbf{0}\} \\ p(\beta) &\propto \phi_{\mathcal{N}_1}(\beta | \mu_\beta, \Sigma_\beta) \mathbb{I}\{\beta > 0\}, \end{aligned}$$

where $\mu.$ and $\Sigma.$ are the hyperparameters, $\mathbb{I}\{\cdot\}$ is the indicator function and $\phi_{\mathcal{N}_d}$ is the d -dimensional Normal density.

The prior distribution of vector $\boldsymbol{\varpi}$ conditional on ν is found by noting that the components ϖ_t are independent and identically distributed from the Inverted Gamma density, which yields:

$$p(\boldsymbol{\varpi} | \nu) = \left(\frac{\nu}{2}\right)^{\frac{T\nu}{2}} \left[\Gamma\left(\frac{\nu}{2}\right)\right]^{-T} \left(\prod_{t=1}^T \varpi_t\right)^{-\frac{\nu}{2}-1} \exp\left[-\frac{1}{2} \sum_{t=1}^T \frac{\nu}{\varpi_t}\right].$$

We follow [Deschamps \(2006\)](#) in the choice of the prior distribution on the degrees of freedom parameter. The distribution is a translated Exponential with parameters $\lambda > 0$ and $\delta \geq 2$:

$$p(\nu) = \lambda \exp[-\lambda(\nu - \delta)] \mathbb{I}\{\nu > \delta\}. \quad (3)$$

Finally, we assume prior independence between the priors.

The recursive nature of the variance equation does not allow for conjugacy between the likelihood function and the prior density. Therefore, we rely on the M-H algorithm to draw samples from the joint posterior distribution. The algorithm is fully described in [Ardia \(2008, Chapter 5\)](#).

3. Illustration

We apply our Bayesian estimation methods to daily observations of the Deutschmark vs British Pound (DEM/GBP) foreign exchange log-returns. The sample period is from January 3, 1985, to December 31, 1991, for a total of 1 974 observations. This data set can be loaded using `data(dem2gbp)`. From this time series, the first 750 observations are used to illustrate the Bayesian approach.

We fit the parsimonious GARCH(1,1) model to the data for this observation window. To that aim, we use the `bayesGARCH` function.

```
R> args(bayesGARCH)
```

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

As a prior distribution for the Bayesian estimation we choose a truncated tri-dimensional Normal distribution with a zero mean vector and a diagonal covariance matrix. The variances are set to 1 000 so we do not introduce tight prior information into our estimation. We run two chains for 1 000 passes each. We emphasize the fact that only positivity constraints are implemented in the MH algorithm; no stationarity conditions are imposed in the simulation procedure.

```
R> set.seed(1)
```

```
R> MCMC <- bayesGARCH(y, control = list(n.chain = 2))
```

```

chain: 1 iteration: 10 parameters: 0.0315 0.2446 0.6619 83.25
chain: 1 iteration: 20 parameters: 0.033 0.1851 0.7234 58.96
...
chain: 2 iteration: 9990 parameters: 0.0464 0.2684 0.6581 4.905
chain: 2 iteration: 10000 parameters: 0.0371 0.2626 0.6602 4.809

```

The sampling algorithm allows to reach very high acceptance rates ranging from 89% for vector α to 95% for β suggesting that the proposal distributions are close to the full conditionals. We discard the first 5'000 draws from the overall MCMC output as a *burn in* period and merge the two chains to get a final sample's length of 10'000.

```
R> smpl <- formSmpl(MCMC, l.bi = 5000)
```

```

n.chain: 2
l.chain: 10000
l.bi: 5000
batch.size: 1
smpl size: 10000

```

The posterior statistics are obtained using the `summary` method of the `coda` package.

```
R> summary(smpl)
```

```

Iterations = 1:10000
Thinning interval = 1
Number of chains = 1
Sample size per chain = 10000

```

1. Empirical mean and standard deviation for each variable, plus standard error of the mean:

	Mean	SD	Naive SE	Time-series SE
alpha0	0.0352	0.0145	0.000145	0.00146
alpha1	0.2450	0.0737	0.000737	0.00716
beta	0.6852	0.0821	0.000821	0.00883
nu	5.7292	1.3740	0.013740	0.14394

2. Quantiles for each variable:

	2.5%	25%	50%	75%	97.5%
alpha0	0.0124	0.0247	0.0334	0.0436	0.0686
alpha1	0.1235	0.1941	0.2375	0.2876	0.4125
beta	0.5251	0.6292	0.6881	0.7432	0.8377
nu	3.6407	4.7270	5.5285	6.5641	8.9994

Marginal distributions are plotted using the `hist` function:

```

> par(mfrow = c(2,2), las = 1)
> smpl <- as.matrix(smpl)
> hist(smpl[,"alpha0"], nclass = 30, col = "grey",
+      border = "white", main = expression(alpha[0]), xlab = "")
> box()
> hist(smpl[,"alpha1"], nclass = 30, col = "grey",
+      border = "white", main = expression(alpha[1]), xlab = "")
> box()
> hist(smpl[,"alpha0"], nclass = 30, col = "grey",
+      border = "white", main = expression(beta), xlab = "")
> box()
> hist(smpl[,"nu"], nclass = 30, col = "grey",
+      border = "white", main = expression(nu), xlab = "")
> box()

```

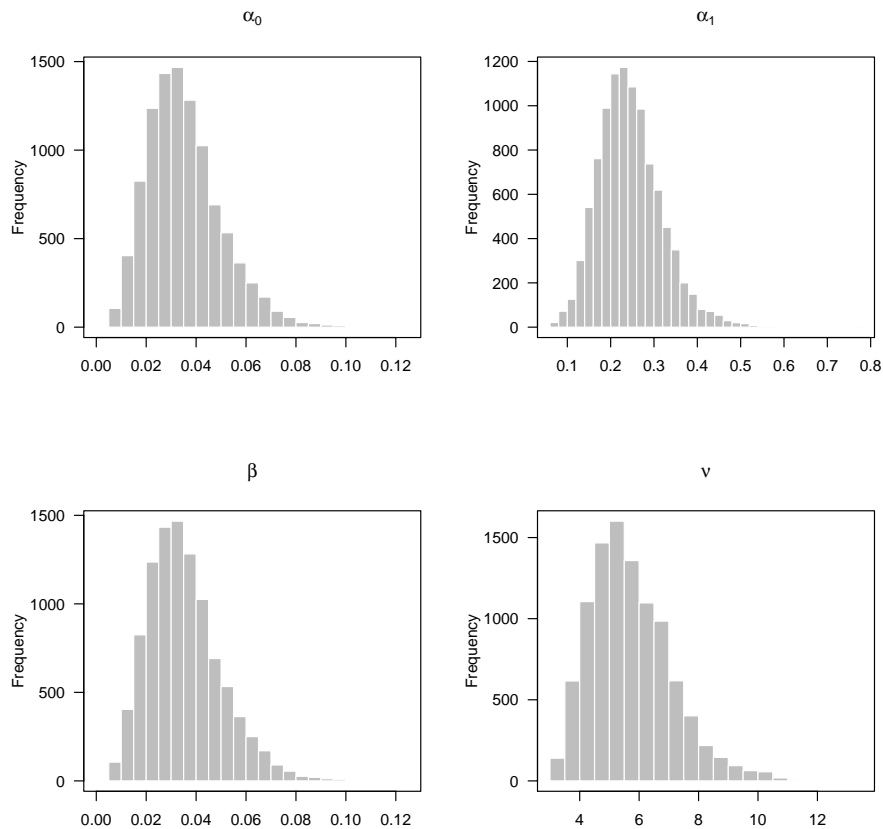


Figure 1: Marginal posterior distributions

3.1. Normal innovations and prior restrictions

The function `addPriorConditions` can be used to impose any type of constraint on the

model parameters ψ . If, e.g., we want to ensure covariance-stationarity, i.e., $\alpha_1 + \beta < 1$. Note that it is important to overwrite the previous function `addPriorConditions` using `assignInNamespace`. Also, we can impose Normality of the innovations in a straightforward manner by setting the hyperparameter $\delta = 500$.

```
R> "addPriorConditions" <- function(psi)
+   psi[2] + psi[3] < 1
R> assignInNamespace("addPriorConditions", addPriorConditions, "bayesGARCH")
R> set.seed(1)
R> MCMC <- bayesGARCH(y, lambda = 100, delta = 500,
+                   control = list(n.chain = 2))
```

```
chain:  1  iteration:  10  parameters:  0.0451 0.1854 0.694 500
chain:  1  iteration:  20  parameters:  0.0492 0.2412 0.6191 500
...
chain:  2  iteration: 9990  parameters:  0.0428 0.1581 0.6917 500
chain:  2  iteration:10000  parameters:  0.0446 0.2928 0.6017 500
```

4. Conclusion

This paper has proposed the Bayesian estimation of GARCH(1,1) model with Student-*t* innovations using the R package **bayesGARCH**. The methodology based on [Ardia \(2008, Chapter 5\)](#) leads to a fast, fully automatic and efficient estimation procedure compared to alternative approaches such as the Griddy-Gibbs sampler. Practitioners who need to run the estimation frequently and/or for a large number of time series should find the procedure helpful. The GARCH(1,1) model has been applied to foreign exchange log-returns time series.

Finally, if you use **bayesGARCH**, please cite the package in publications. Use:

```
> citation("bayesGARCH")
```

Computational details

The results in this paper were obtained using R 2.8.1 ([R Development Core Team 2008](#)) with the packages **bayesGARCH** 1.00-01 ([Ardia 2007](#)) and **coda** 0.13-2 ([Plummer, Best, Cowles, and Vines 2008](#)). R itself and all packages used are available from CRAN at <http://CRAN.R-project.org/>. Computations were performed on a Genuine Intel® dual core CPU T2400 1.83Ghz processor. Code outputs were obtained using `options(digits=4, max.print=40)`.

Acknowledgements

This version has been written while the author was visiting CORE, Université catholique de Louvain, Belgium. The author sincerely acknowledges the hospitality of Luc Bauwens and is grateful to the Swiss National Science Foundation (under grant #FN PB FR1-121441) for financial support. Any remaining errors or shortcomings are the author's responsibility.

References

- Ardia D (2007). ‘*bayesGARCH*’: *Bayesian Estimation of the GARCH(1,1) Model with Student-t Innovations in R*. URL <http://CRAN.R-project.org/package=AdMit>.
- Ardia D (2008). *Financial Risk Management with Bayesian Estimation of GARCH Models: Theory and Applications*, volume 612 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, Germany. ISBN 978-3-540-78656-6. doi:10.1007/978-3-540-78657-3.
- Deschamps PJ (2006). “A Flexible Prior Distribution for Markov Switching Autoregressions with Student-t Errors.” *Journal of Econometrics*, **133**(1), 153–190. doi:10.1016/j.jeconom.2005.03.012.
- Engle RF (1982). “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation.” *Econometrica*, **50**(4), 987–1008.
- Plummer M, Best N, Cowles K, Vines K (2008). *coda: Output Analysis and Diagnostics for MCMC in R*. R package version 0.13-3, URL <http://CRAN.R-project.org/package=coda>.
- R Development Core Team (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org/>.

Affiliation:

David Ardia
Department of Quantitative Economics
University of Fribourg
CH 1700 Fribourg
Switzerland
E-mail: david.ardias@unifr.ch
URL: <http://perso.unifr.ch/david.ardias/>