WIRTSCHAFTS
UNIVERSITÄT
WIEN VIENNA
UNIVERSITY OF
ECONOMICS
AND BUSINESS

Dealing with stochastic volatility in time series using the R package stochvol

Gregor Kastner

R/Finance 2014, May 16–17

Part I

A pithy introduction to SV

3140 EUR/USD daily exchange rates

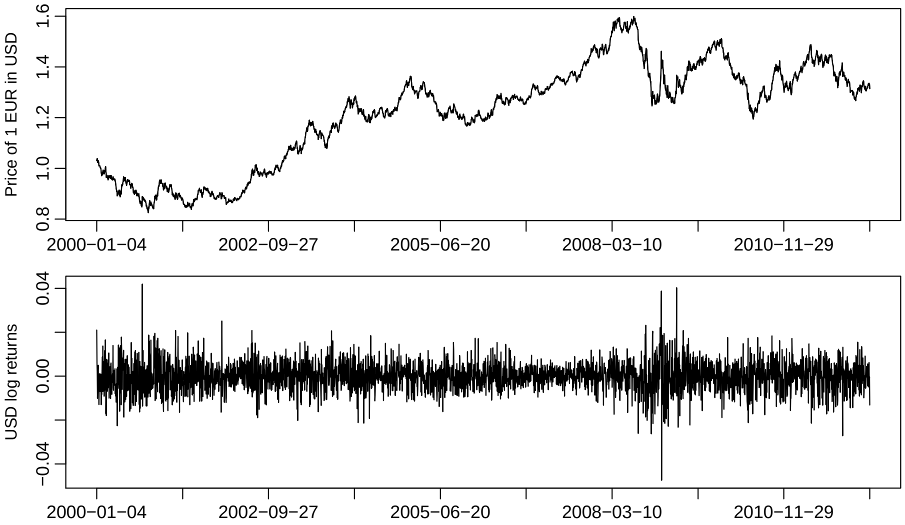

Autocorrelation function

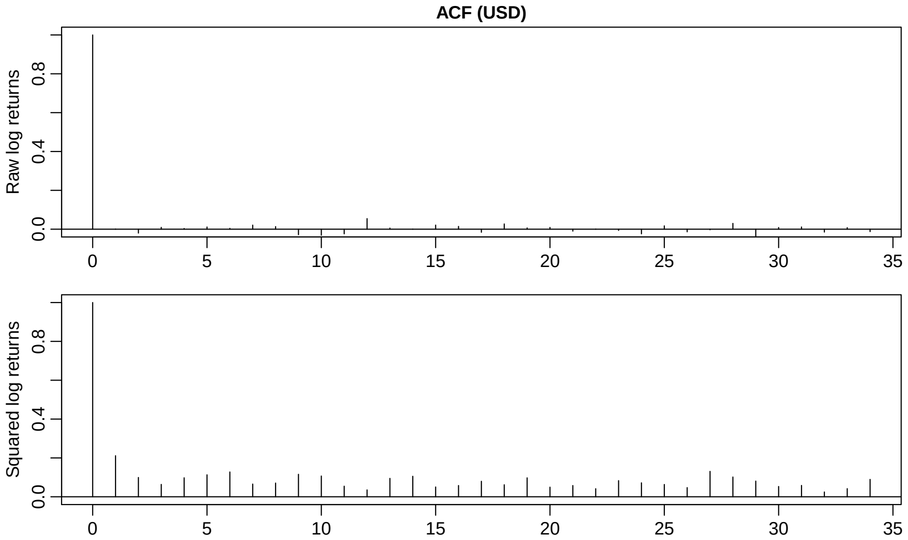

GARCH vs. SV

GARCH(1,1)

$$\begin{aligned} \tilde{y}_t &= h_t \varepsilon_t, & \varepsilon_t \sim N(0,1), \\ h_t^2 &= \alpha_0 + \alpha_1 \tilde{y}_{t-1}^2 + \beta h_{t-1}^2, \end{aligned}$$

with h_0, α_0, α_1 and β unknown.

GARCH vs. SV

GARCH(1,1)

$$\begin{aligned} \tilde{y}_t &= h_t \varepsilon_t, & \varepsilon_t \sim N(0,1), \\ h_t^2 &= \alpha_0 + \alpha_1 \tilde{y}_{t-1}^2 + \beta h_{t-1}^2, & \end{aligned}$$

with h_0, α_0, α_1 and β unknown.

SV

$$\begin{aligned} \tilde{y}_t &= h_t \varepsilon_t, & \varepsilon_t \sim N(0,1), \\ \log h_t^2 &= \mu + \phi(\log h_{t-1}^2 - \mu) + \sigma \eta_t, & \eta_t \sim N(0,1), \end{aligned}$$

with $\mathbf{h} = (h_0, \ldots, h_T)$, μ, ϕ and σ unknown.

Stochastic volatility models . . .

- have separate error terms for the mean and volatility equations (state space models).
- arise as discrete approximations to various diffusion processes in continuous-time asset pricing.
- shares many "nice" properties with (X-)GARCH . . .
- . . . while still showing important differences, e.g. a more realistic ACF of squared returns.
- shows less/no remaining nonlinear dependencies in residuals (Jacquier, Polson, and Rossi, 1994).
- have a likelihood that is difficult to evaluate.

Stochastic volatility models ...

- have separate error terms for the mean and volatility equations (state space models).
- arise as discrete approximations to various diffusion processes in continuous-time asset pricing.
- shares many "nice" properties with (X-)GARCH ...
- ... while still showing important differences, e.g. a more realistic ACF of squared returns.
- shows less/no remaining nonlinear dependencies in residuals (Jacquier, Polson, and Rossi, 1994).
- have a likelihood that is difficult to evaluate. Solutions:
  - MM (no score function available, relatively inefficient)
  - approximate linear filtering methods for a QML estimator
  - other methods for approximating the integral used in evaluating the likelihood (revival?)
  - Bayesian methods (mainly MCMC)

Model

"Centered" version (e.g. Jacquier, Polson, and Rossi, 1994):

$$
\begin{aligned}
y_t &= e^{h_t/2}\,\varepsilon_t, & \varepsilon_t &\sim N(0,1), \\
h_t &= \mu + \phi(h_{t-1} - \mu) + \sigma\eta_t, & \eta_t &\sim N(0,1),
\end{aligned}
$$

with $\mathbf{h} = (h_0, \ldots, h_T)$, μ, ϕ and σ unknown.

Model

"Centered" version (e.g. Jacquier, Polson, and Rossi, 1994):

$$\begin{aligned} y_t &= e^{h_t/2}\,\varepsilon_t, & \varepsilon_t &\sim N(0,1), \\ h_t &= \mu + \phi(h_{t-1}-\mu) + \sigma\eta_t, & \eta_t &\sim N(0,1), \end{aligned}$$

with $\mathbf{h} = (h_0, \ldots, h_T)$, μ, ϕ and σ unknown. Priors:

- $\mu \sim N(b_\mu, B_\mu)$
- $(\phi+1)/2 \sim \mathcal{B}(a_0, b_0)$ as in Kim, Shephard, and Chib (1998), implying

$$p(\phi) = \frac{1}{2B(a_0, b_0)} \left(\frac{1+\phi}{2}\right)^{a_0-1} \left(\frac{1-\phi}{2}\right)^{b_0-1}$$

- $\sigma^2 \sim B_\sigma \cdot \chi_1^2 = \mathcal{G}\left(\frac{1}{2}, \frac{1}{2B_\sigma}\right)$
- $h_0|\mu,\phi,\sigma \sim N\left(\mu, \sigma^2/(1-\phi^2)\right)$

Estimation I

Kastner and Frühwirth-Schnatter (forthcoming) provide an MCMC algorithm for efficient univariate SV estimation.

Estimation I

Kastner and Frühwirth-Schnatter (forthcoming) provide an MCMC algorithm for efficient univariate SV estimation. Key features:

- AWOL: Sample the latent variables jointly, All WithOut a Loop: Rue (2001), McCausland, Miller, and Pelletier (2011):
  - No FFBS needed
  - No need to invert (band diagonal) precision matrix of latents
  - Fast (band back-substitution)
  - Implemented in C, Matlab, R, ...

Estimation I

Kastner and Frühwirth-Schnatter (forthcoming) provide an MCMC algorithm for efficient univariate SV estimation. Key features:

- AWOL: Sample the latent variables jointly, All WithOut a Loop: Rue (2001), McCausland, Miller, and Pelletier (2011):
  - No FFBS needed
  - No need to invert (band diagonal) precision matrix of latents
  - Fast (band back-substitution)
  - Implemented in C, Matlab, R, ...
- ASIS: Ancillarity-Sufficiency Interweaving Strategy (Yu and Meng, 2011). Sample the "critical but cheap" parameters μ, σ (and ϕ), twice: Once centered, and once noncentered, i.e. based on:

$$
\begin{aligned}
y_t &= e^{(\mu+\sigma\tilde{h}_t)/2}\,\varepsilon_t, & \varepsilon_t &\sim N(0,1), \\
\tilde{h}_t &= \phi\tilde{h}_{t-1} + \eta_t, & \eta_t &\sim N(0,1).
\end{aligned}
$$

Estimation II

$\log(\varepsilon_t^2)|r_t \sim N\left(m_{r_t}, s_{r_t}^2\right)$ for $r_t \in \{1, \ldots, 10\}$ (Omori et al., 2007):

1 Sample $\mathbf{h} = (h_0, \ldots, h_T)$ AWOL:
   - $\mathbf{h} = (\mathbf{L}')^{-1}(\mathbf{L}^{-1}\mathbf{c} + \boldsymbol{\varepsilon})$ with $\boldsymbol{\Omega} = \mathbf{L}\mathbf{L}'$ and $\boldsymbol{\varepsilon} \sim N_T(\mathbf{0}, \mathbf{I}_T)$

2 Sample μ, ϕ, σ (C):
   - MH-step with independence proposal for $(\mu, \phi)|\mathbf{h}, \sigma^2$
   - MH-step with independence proposal for $\sigma^2|\mathbf{h}, \mu, \phi$

2* Calculate $\tilde{\mathbf{h}} = \frac{\mathbf{h}-\mu}{\sigma}$

2** Sample μ, ϕ, σ (NC):
   - MH-step with independence proposal for $\phi|\tilde{\mathbf{h}}$
   - Gibbs-step for $(\mu, \sigma)|\mathbf{y}, \tilde{\mathbf{h}}, \mathbf{r}$ with possible sign switching for σ

2*** Calculate $\mathbf{h} = \mu + \sigma\tilde{\mathbf{h}}$ ("bookkeeping")

3 Draw indicators r_t from easily available posterior

Inefficiency factors for $p(\mu|\mathbf{y})$: 500×100k draws, $T = 5$k

$p(\mu\|\mathbf{y})$	σ_{true} \ ϕ_{true}	0	0.5	0.8	0.9	0.95	0.96	0.97	0.98	0.99
C 2.31s	0.1	641	223	64	18	6	5	3	2	3
	0.2	253	89	22	9	3	2	2	2	3
	0.3	113	47	15	5	2	2	2	2	4
	0.4	72	33	11	3	2	2	2	2	5
	0.5	52	26	8	2	2	2	2	2	5
NC 2.34s	0.1	9	10	12	13	21	30	50	113	487
	0.2	24	21	15	22	70	108	190	419	1729
	0.3	23	17	17	41	148	234	407	922	3743
	0.4	18	16	23	70	265	412	726	1707	6790
	0.5	17	17	32	105	411	660	1149	2534	9421
GIS-C 2.36s	0.1	9	9	11	8	4	3	3	2	3
	0.2	23	20	11	5	2	2	2	2	3
	0.3	22	15	9	4	2	2	2	2	3
	0.4	17	13	7	3	2	2	2	2	4
	0.5	15	12	5	2	2	2	2	2	4

Inefficiency factors for $p(\sigma|\mathbf{y})$: 500 $\times$ 100k draws, $T = 5$k

$p(\sigma\|\mathbf{y})$	σ_{true} \ ϕ_{true}	0	0.5	0.8	0.9	0.95	0.96	0.97	0.98	0.99
C 2.31s	0.1	5440	4899	3608	1845	768	604	431	300	194
	0.2	3001	1813	779	374	181	151	124	100	80
	0.3	663	490	316	155	90	79	68	59	50
	0.4	238	260	163	91	58	54	48	43	38
	0.5	135	166	105	63	44	41	37	34	31
NC 2.34s	0.1	57	64	91	130	90	83	73	71	87
	0.2	99	95	121	87	70	70	75	89	137
	0.3	61	77	94	70	71	77	89	114	189
	0.4	38	75	75	66	79	90	106	144	257
	0.5	31	68	67	67	90	103	126	173	316
GIS-C 2.36s	0.1	56	64	89	122	82	75	64	58	61
	0.2	97	93	114	76	53	50	48	48	51
	0.3	58	72	82	53	42	41	40	39	40
	0.4	35	69	58	42	35	35	34	33	33
	0.5	28	59	47	35	31	30	29	29	28

A quote from Bos (2012)

While there are literally thousands of applications of GARCH, for SV, this number is far lower. Two reasons for this relative lack of applied work using SV are apparent. First, there are as of yet no standard packages for estimating SV models, whereas for GARCH, most statistical packages have a wealth of options for incorporating GARCH effects. [...]

Part II
The stochvol package (Kastner, 2014)

Running the sampler

```
R> library(stochvol)
R> data(exrates)
R> ret <- logret(exrates$USD, demean = TRUE)
R> res <- svsample(ret, priormu    = c(-10, 1),
+                       priorphi   = c(20, 1.1),
+                       priorsigma = .1)
```

```
Calling GIS_C MCMC sampler with 11000 iter. Series length T = 3139.

  0% [++++++++++++++++++++++++++++++++++++++++++++++++++++] 100%

Timing (elapsed): 17.935 seconds.
613 iterations per second.

Converting results to coda objects... Done!
Summarizing posterior draws... Done!
```

Printing and plotting the resulting object

```
R> summary(res, showlatent = FALSE)
```

```
Summary of 10000 MCMC draws after a burn-in of 1000.

Prior distributions:
mu          ~ Normal(mean = -10, sd = 1)
(phi+1)/2 ~ Beta(a0 = 20, b0 = 1.1)
sigma^2     ~ 0.1 * Chisq(df = 1)

Posterior draws of parameters (thinning = 1):
                mean      sd       5%      50%      95%  ESS
mu          -10.1308 0.22806 -10.4661 -10.1345 -9.7746 4537
phi           0.9936 0.00282   0.9886   0.9938  0.9977  386
sigma         0.0654 0.01005   0.0510   0.0646  0.0827  142
exp(mu/2)     0.0064 0.00076   0.0053   0.0063  0.0075 4537
sigma^2       0.0044 0.00139   0.0026   0.0042  0.0068  142
```

```
R> plot(res)
```

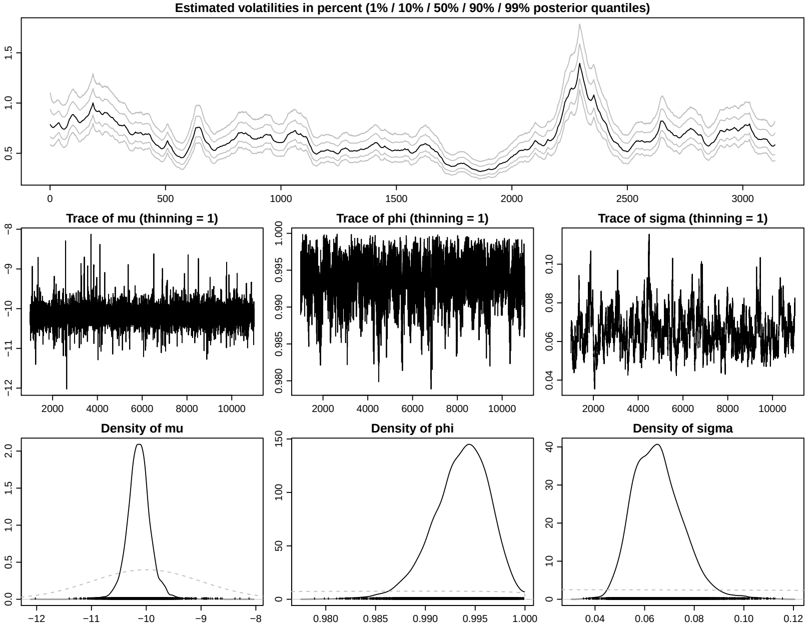

Estimated volatilities in percent (1% / 10% / 50% / 90% / 99% posterior quantiles)
Trace of mu (thinning = 1)
Trace of phi (thinning = 1)
Trace of sigma (thinning = 1)
Density of mu
Density of phi
Density of sigma

```
R> plot(resid(res), ret)
```

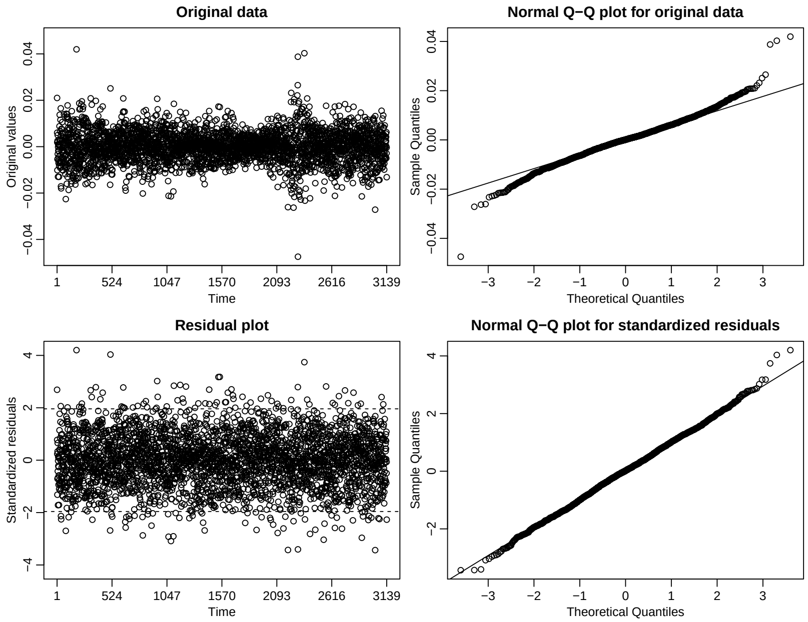

```
R> volplot(res, forecast = 100, dates = exrates$date[-1])
```

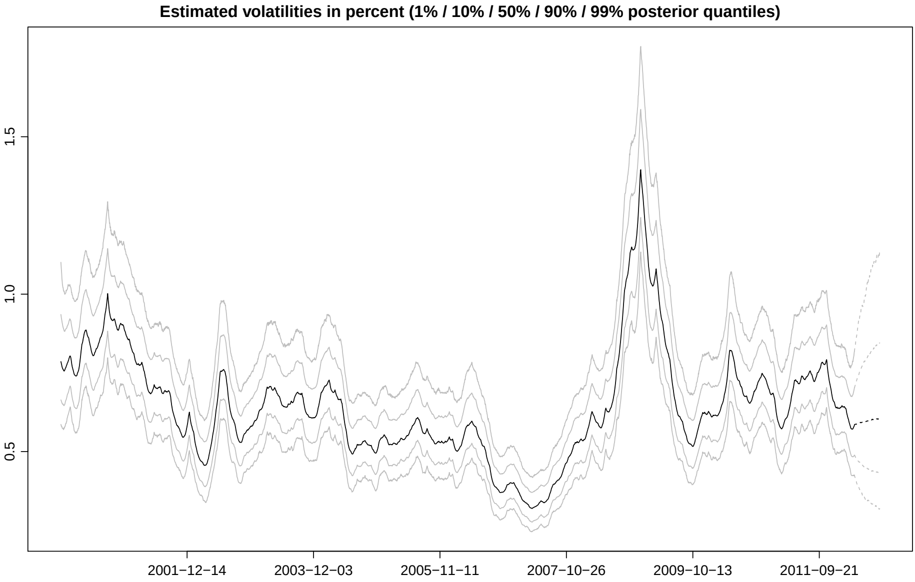

Part III

Using stochvol within other samplers

Consider a Bayesian normal linear model with T observations and $k = p - 1$ predictors,

$$\mathbf{y} | \boldsymbol{\beta}, \boldsymbol{\Sigma} \sim N(\mathbf{X}\boldsymbol{\beta}, \boldsymbol{\Sigma}).$$

Two specifications of errors:

- Homoskedastic: $\boldsymbol{\Sigma} \equiv \sigma_\epsilon^2 \mathbf{I}$
- Heteroskedastic: $\boldsymbol{\Sigma} \equiv \mathrm{Diag}\left(e^{h_1}, \ldots, e^{h_T}\right)$

Bayesian regression: Concept

Consider a Bayesian normal linear model with T observations and $k = p - 1$ predictors,

$$\mathbf{y}|\boldsymbol{\beta}, \boldsymbol{\Sigma} \sim N(\mathbf{X}\boldsymbol{\beta}, \boldsymbol{\Sigma}).$$

Two specifications of errors:

- Homoskedastic: $\boldsymbol{\Sigma} \equiv \sigma_\epsilon^2 \mathbf{I}$
- Heteroskedastic: $\boldsymbol{\Sigma} \equiv \text{Diag}\left(e^{h_1}, \dots, e^{h_T}\right)$

Gibbs sampler for drawing from the posterior:

- Homoskedastic: Using conjugate priors, a simple Gibbs sampler is given through drawing in turn from the full conditional distributions

$$\begin{aligned} \boldsymbol{\beta}|\mathbf{y}, \boldsymbol{\Sigma} &\sim N(\mathbf{b}_T, \mathbf{B}_T) && (1) \\ \sigma_\epsilon^2|\mathbf{y}, \boldsymbol{\beta} &\sim \mathcal{G}^{-1}(c_T, C_T) && (2) \end{aligned}$$

- Heteroskedastic: Simply replace (2) by a call to `svsample`!

Daily price of 1 EUR in USD from January 3, 2000 until April 4, 2012, denoted by $\mathbf{p} = (p_1, p_2, \ldots, p_T)'$. Let

$$\mathbf{y} = \begin{pmatrix} p_2 \\ p_3 \\ \vdots \\ p_T \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} 1 & p_1 \\ 1 & p_2 \\ \vdots & \vdots \\ 1 & p_{T-1} \end{pmatrix}.$$

Clearly, we expect the posterior distribution of β to spread around $(0, 1)'$ which corresponds to a random walk.

Bayesian regression: Baby example

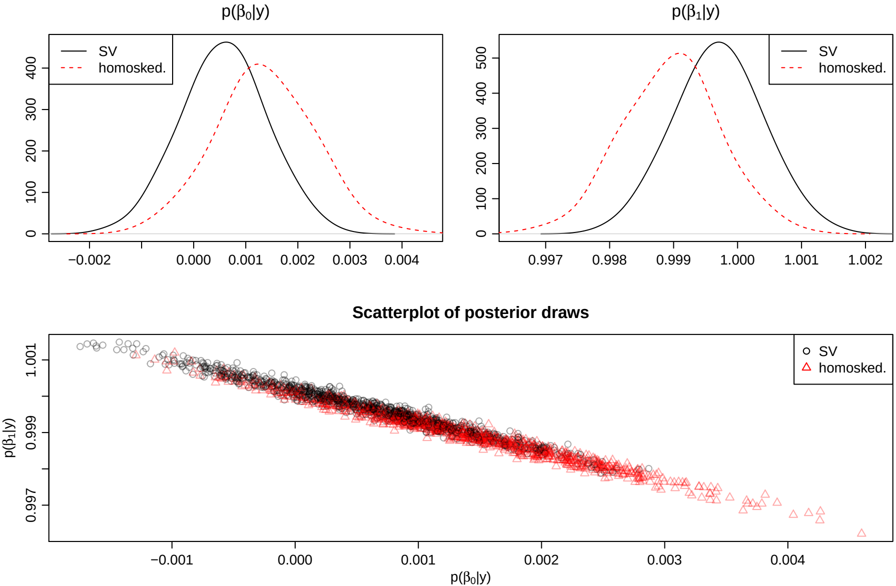

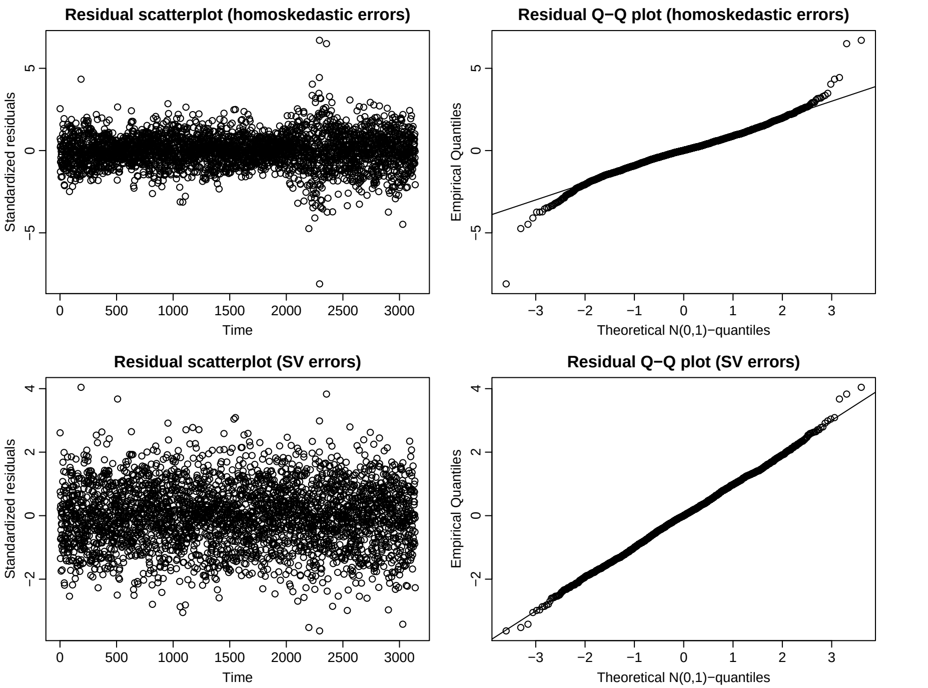
Residual scatterplot (homoskedastic errors)
Residual Q−Q plot (homoskedastic errors)
Residual scatterplot (SV errors)
Residual Q−Q plot (SV errors)
Standardized residuals
Empirical Quantiles
Time
Theoretical N(0,1)−quantiles

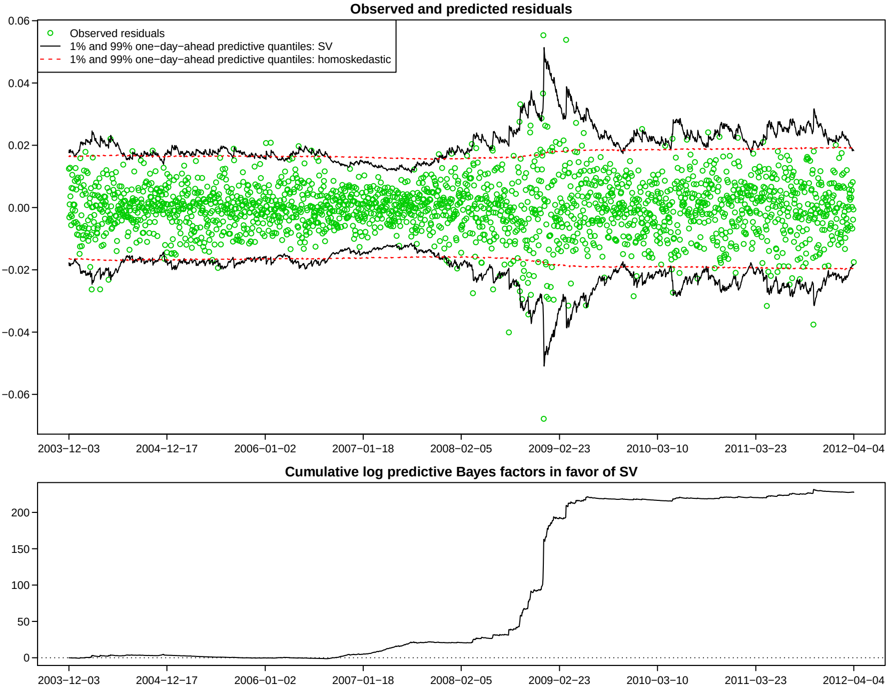
Observed and predicted residuals
Observed residuals
1% and 99% one−day−ahead predictive quantiles: SV
1% and 99% one−day−ahead predictive quantiles: homoskedastic
0.06
0.04
0.02
0.00
−0.02
−0.04
−0.06
2003−12−03
2004−12−17
2006−01−02
2007−01−18
2008−02−05
2009−02−23
2010−03−10
2011−03−23
2012−04−04
Cumulative log predictive Bayes factors in favor of SV
200
150
100
50
0

Factor SV: Motivation

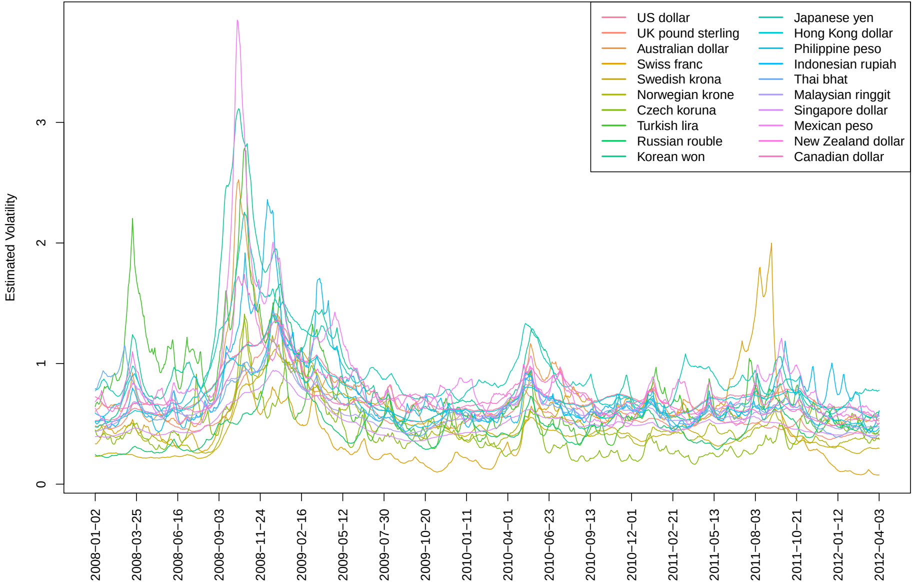

Factor SV (Chib, Nardari, and Shephard, 2006)

$$
\begin{aligned}
\tilde{\mathbf{y}}_t &= \boldsymbol{\Lambda}\mathbf{f}_t + \boldsymbol{\Sigma}_t^{1/2}\varepsilon_t, \qquad \varepsilon_t \sim N_m(\mathbf{0}, \mathbf{I}_m), \\
\mathbf{f}_t &= \mathbf{V}_t^{1/2}\mathbf{u}_t, \qquad \mathbf{u}_t \sim N_r(\mathbf{0}, \mathbf{I}_r),
\end{aligned}
$$

$$
\begin{bmatrix} \boldsymbol{\Sigma}_t & \mathbf{O} \\ \mathbf{O} & \mathbf{V}_t \end{bmatrix} = \text{Diag}(\exp(h_{1,t}), \ldots, \exp(h_{m+r,t})),
$$

where $\tilde{\mathbf{y}}_t = (\tilde{y}_{1t}, \ldots, \tilde{y}_{mt})'$ is the m-variate vector of (potentially demeaned) log-returns at time $t = 1, \ldots, T$, and $\boldsymbol{\Lambda}$ is an unknown $m \times r$ factor loading matrix. $\mathbf{f}_t$, $\mathbf{f}_s$, ε_t, and ε_s are assumed to be pairwise independent for all t and s. Latent factors and the idiosyncratic shocks follow different univariate SV processes, i.e.

$$
h_{it} = \mu_i + \phi_i(h_{i,t-1} - \mu_i) + \sigma_i\eta_{it}, \qquad \eta_{it} \sim N(0, 1).
$$

Factor SV: Sampling scheme

Choose $T(m+2r) + mr + 4m + 3r$, in our application 81 763, starting values for $\boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\sigma}, \mathbf{h}, \mathbf{f}, \boldsymbol{\Lambda}$ and repeat:

a) Perform $m + r$ univariate SV updates for $\mathbf{h}_i, (\mu_i), \phi_i, \sigma_i$ by using **stochvol**.
b) Sample the factor loadings, constituting m independent r-variate regression problems with T observations.
c) Sample the latent factors, constituting T independent r-variate regression problems with m observations.

Note that in principle these steps can be straightforwardly parallelized.

Median posterior factor loadings Λ (sign identifier: JPY)

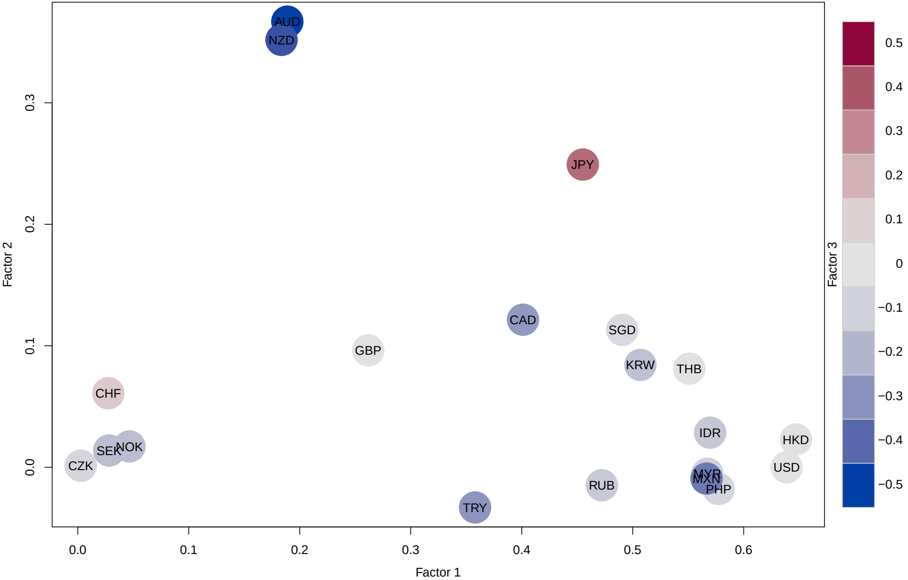

Median posterior communalities $1 - \frac{\Sigma_{ii,t}}{\mathbb{V}(\mathbf{y}_{i,t})}$

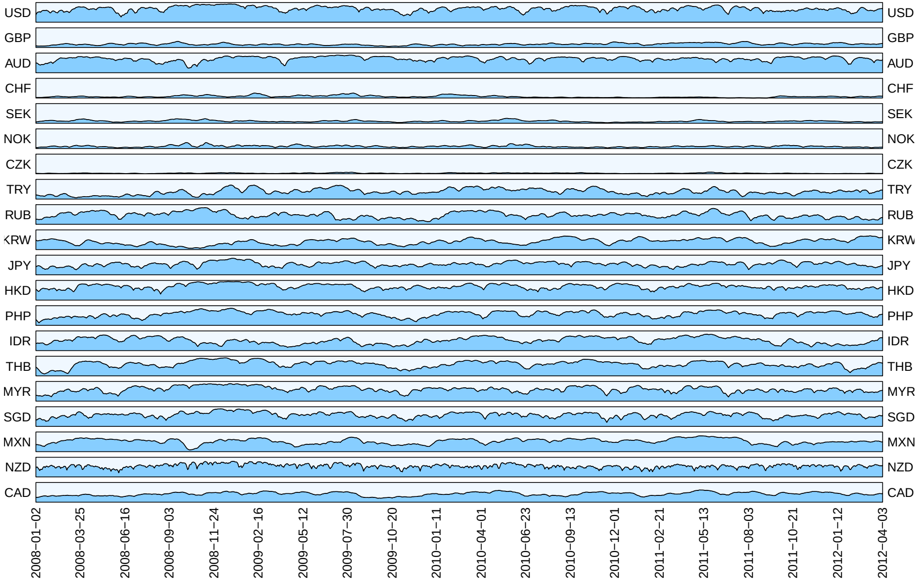

Three posterior median correlation matrices

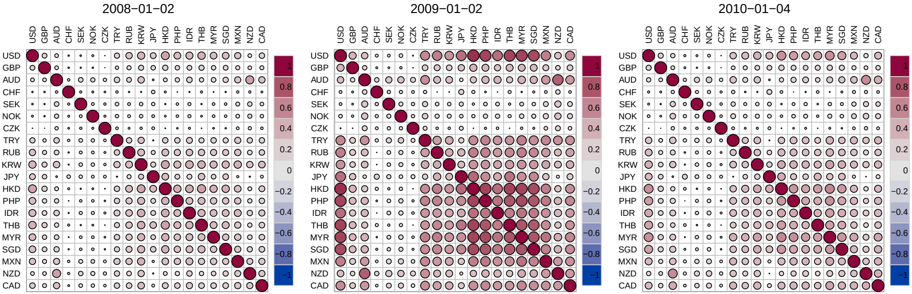

(as implied by the covariance matrix given through $\mathbf{\Lambda V}_t\mathbf{\Lambda}' + \mathbf{\Sigma}_t$)

Some more details in Kastner, Frühwirth-Schnatter, and Lopes, 2014.

Factor SV with hierarchical shrinkage

Regularization for the factor loadings $\mathbf{\Lambda}$:

- Ridge: L^2-penalty (Gaussian prior) – little shrinkage
- LASSO: L^1-penalty (Laplace prior) – (much) more shrinkage
- ...

Factor SV with hierarchical shrinkage

Regularization for the factor loadings $\boldsymbol{\Lambda}$:

- Ridge: L^2-penalty (Gaussian prior) – little shrinkage
- LASSO: L^1-penalty (Laplace prior) – (much) more shrinkage
- ...

Possible ways out:

- Elastic net (combination of Ridge and LASSO)
- Normal-Gamma prior (Griffin and Brown, 2010):

$$\Lambda_{ij}|\psi_{ij} \sim N\left(0, \psi_{ij}\right), \quad \psi_{ij} \sim \mathcal{G}\left(\lambda_i, 1/(2\gamma_i^2)\right),$$

where λ_i is a fixed structural parameter, $2\lambda_i\gamma_i^2 \sim \mathcal{G}^{-1}\left(2, M_i\right)$.

Note that the variance of Λ_{ij} is $2\lambda_i\gamma_i^2$ (which itself has expectation M_i) and the excess kurtosis of Λ_{ij} is $3/\lambda_i$.

Sampling the factor loadings under the NG-prior

For each $i = 1, \ldots, m$ do the following:

b*) Sample γ_i^{-2} from its full conditional distribution, which is $\mathcal{G}\left(2 + n\lambda_i, M_i/2\lambda_i + \frac{1}{2}\sum_{j=1}^{n} \psi_{ij}\right)$.

b**) Update ψ_{ij} (independently for each j) from the Generalized Inverse Gaussian distribution $\mathrm{GIG}(\lambda_i - \frac{1}{2}, \gamma_i^{-2}, \Lambda_{ij}^2)$, where $\mathrm{GIG}(m, c, d)$ has a density proportional to

$$x^{m-1} \exp\left\{-\frac{1}{2}(cx + d/x)\right\}.$$

This can be achieved e.g. through the R package **GIGrvg** (Leydold and Hörmann, 2014).

b***) Draw $\boldsymbol{\Lambda}_{i\cdot}'|\mathbf{f}, \mathbf{y}_i, \mathbf{h}_i, \boldsymbol{\Psi}_{i\cdot} \sim N_r\left(\mathbf{b}_{iT}, \mathbf{B}_{iT}\right)$.

S&P 500 application

- Seven-factor SV model to stock prices listed in the S&P 500 index.
- Only firms which have been listed from November 1994 onwards, resulting in $m = 300$ stock prices on 5001 days, ranging from 11/1/1994 to 12/31/2013.
- Data was obtained from Bloomberg Terminal in January 2014.
- Investigate $T = 5000$ demeaned percentage log-returns.

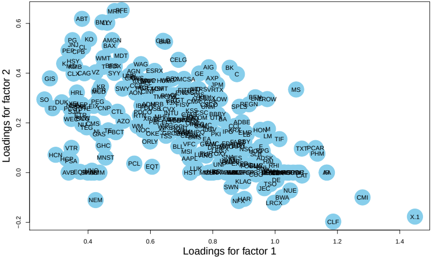

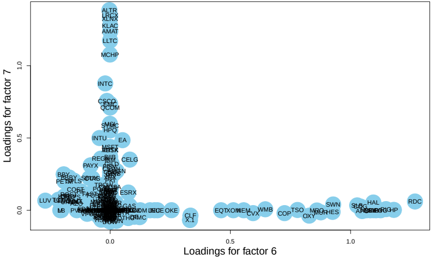

Median correlation matrix on 11/14/2007

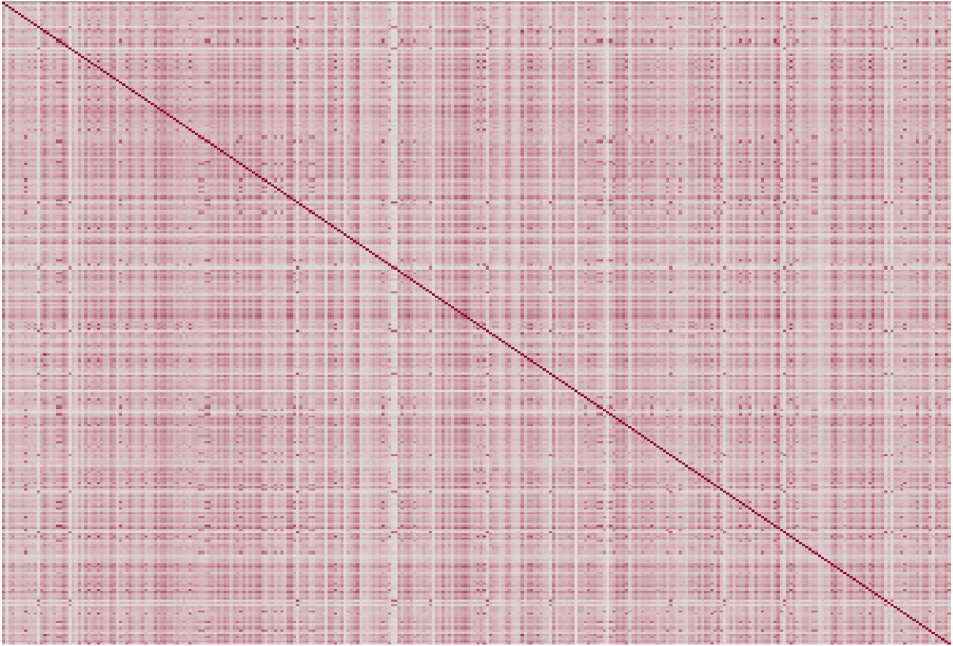

Median correlation matrix on 10/29/2008

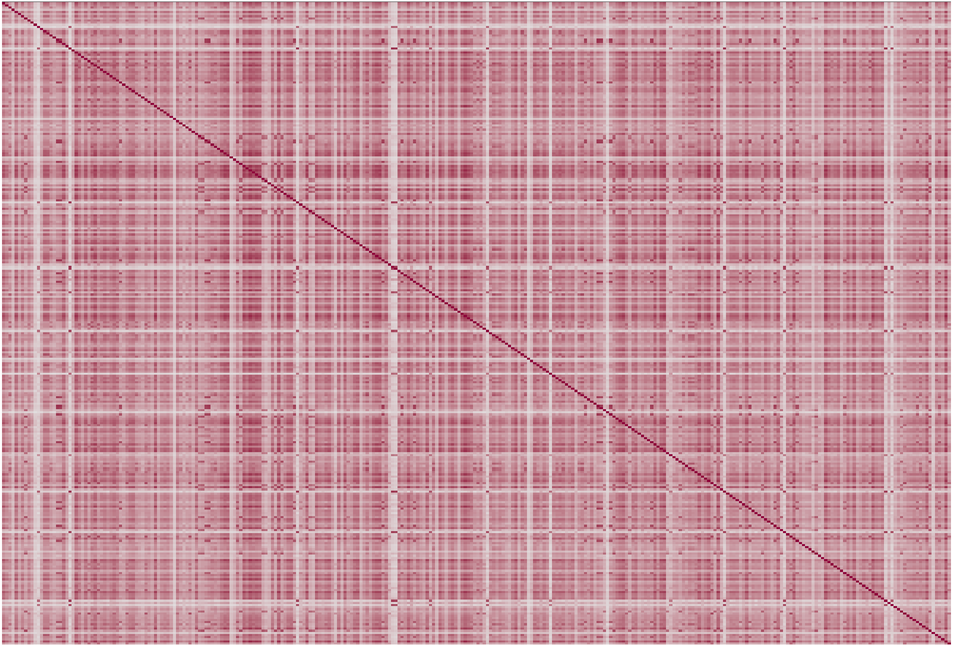

Median correlation matrix on 8/5/2009

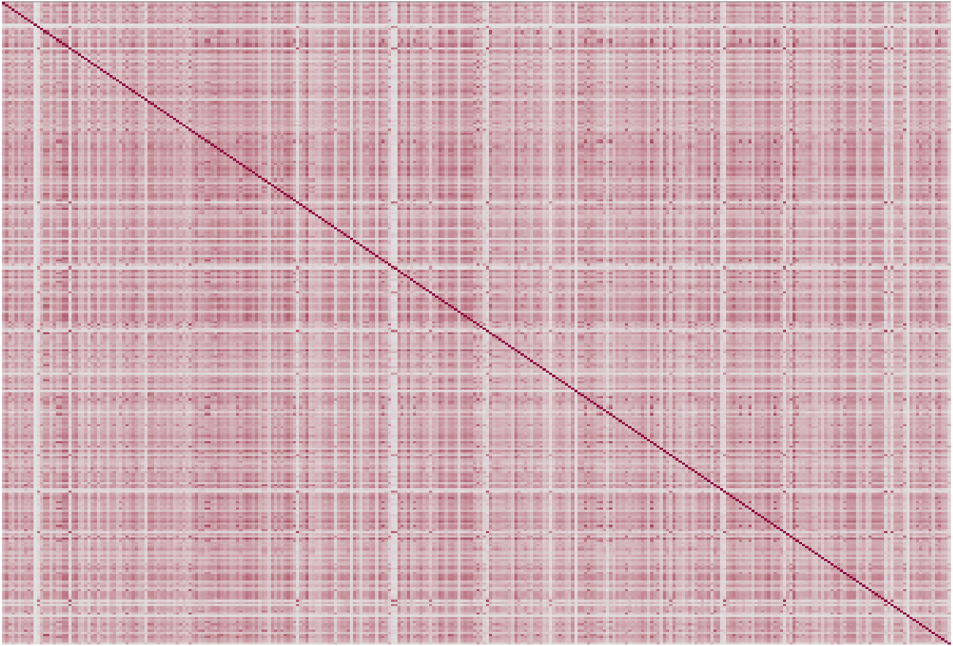

Median correlation matrix on 5/12/2010

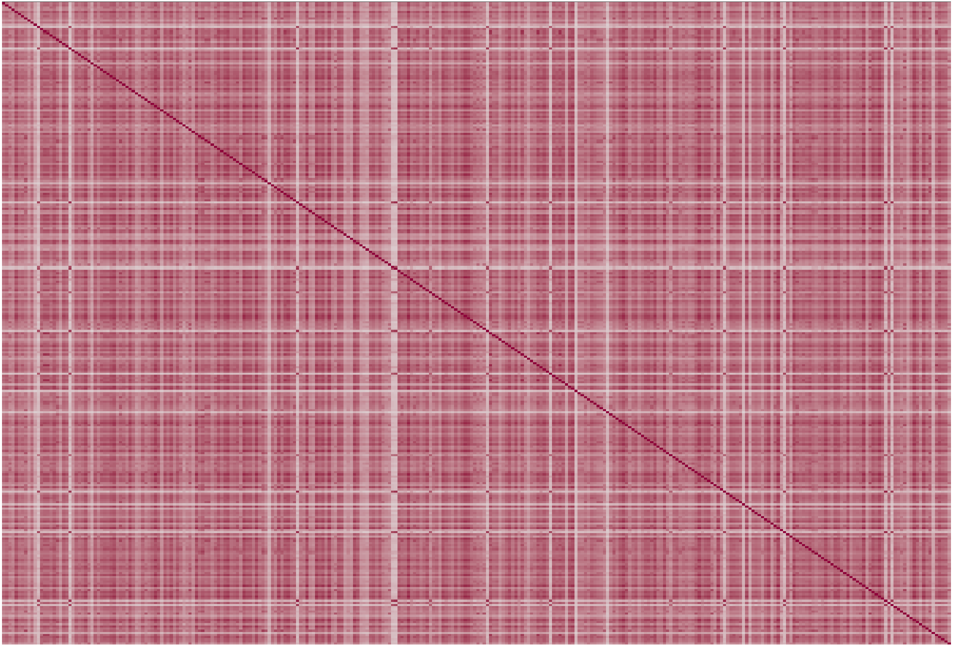

Median correlation matrix on 4/27/2011

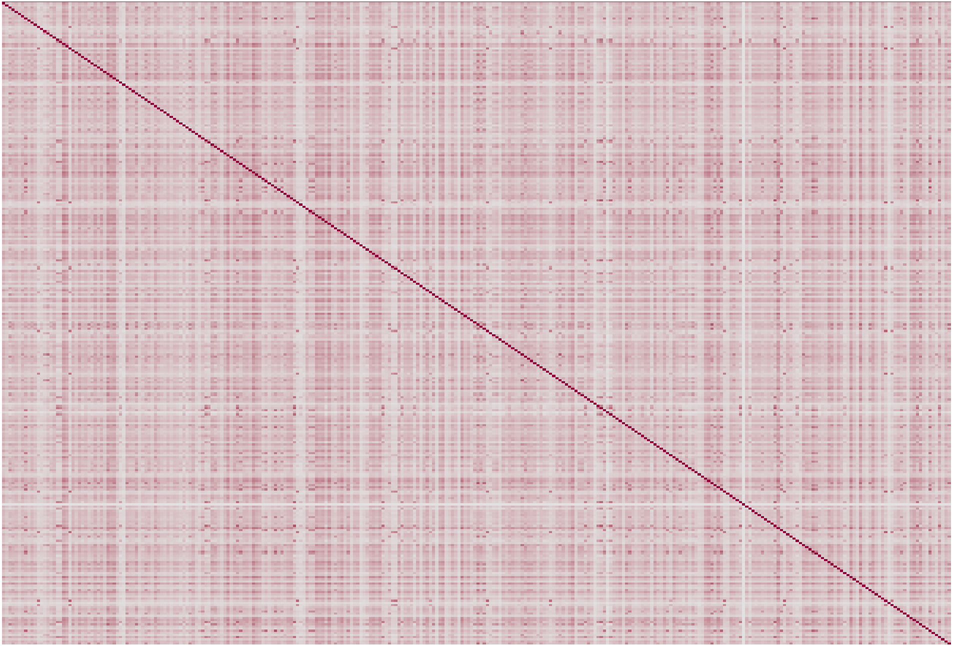

Work in progress

- Improve computation time by fully employing a compiled programming language. Currently only the univariate SV part is implemented in C/C++ via **Rcpp** (Eddelbuettel and François, 2011).
- NEW: Promising preliminary results using **RcppArmadillo** (Eddelbuettel and Sanderson, 2014) for the entire sampler (speedup 50+fold).
- Improve computation time by (possibly massively) parallel implementation. Currently singleCPU/singlecore for r small and singleCPU/multicore for r large (when linked against BLAS).
- Further increase dimensionality (current state: $m \approx 1000$, $r \approx 20$).
- Further increase simulation efficiency by extending the interweaving strategy for sampling the factor loadings. Note e.g. that $\mu_i = 2 \log \Lambda_{ii}$ for $i > m$.
- Finalize **factorstochvol** and put it on CRAN...

References I

Bos, C. S. (2012). "Relating stochastic volatility estimation methods". In: **Handbook of Volatility Models and Their Applications**. Ed. by L. Bauwens, C. Hafner, and S. Laurent. Wiley, pp. 147–174. DOI: 10.1002/9781118272039.ch6 (cit. on p. 17).

Chib, S., F. Nardari, and N. Shephard (2006). "Analysis of high dimensional multivariate stochastic volatility models". **Journal of Econometrics** 134, pp. 341–371. DOI: 10.1016/j.jeconom.2005.06.026 (cit. on p. 32).

Eddelbuettel, D. and R. François (2011). "Rcpp: Seamless R and C++ Integration". **Journal of Statistical Software** 40, pp. 1–18. URL: http://www.jstatsoft.org/v40/i08/ (cit. on p. 46).

Eddelbuettel, D. and C. Sanderson (2014). "RcppArmadillo: Accelerating R with high-performance C++ linear algebra". **Computational Statistics and Data Analysis** 71, pp. 1054–1063. DOI: 10.1016/j.csda.2013.02.005 (cit. on p. 46).

Griffin, J. E. and P. J. Brown (2010). "Inference with Normal-Gamma prior distributions in regression problems". **Bayesian Analysis** 5, pp. 171–188. DOI: 10.1214/10-BA507 (cit. on pp. 37, 38).

Jacquier, E., N. G. Polson, and P. E. Rossi (1994). "Bayesian analysis of stochastic volatility models". **Journal of Business & Economic Statistics** 12, pp. 371–389. DOI: 10.1080/07350015.1994.10524553 (cit. on pp. 7–10).

References II

Kastner, G. (2014). **stochvol: Efficient Bayesian inference for stochastic volatility (SV) models**. R package version 0.8-2. URL: http://CRAN.R-project.org/package=stochvol (cit. on pp. 18–23).

Kastner, G. and S. Frühwirth-Schnatter (forthcoming). "Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models". **Computational Statistics and Data Analysis**. DOI: 10.1016/j.csda.2013.01.002 (cit. on pp. 11–13).

Kastner, G., S. Frühwirth-Schnatter, and H. F. Lopes (2014). "Analysis of exchange rates via multivariate Bayesian factor stochastic volatility models". In: **The Contribution of Young Researchers to Bayesian Statistics – Proceedings of BAYSM2013**. Ed. by E. Lanzarone and I. Francesca. Vol. 63. Springer Proceedings in Mathematics & Statistics. Springer, pp. 181–186. DOI: 10.1007/978-3-319-02084-6_35 (cit. on p. 36).

Kim, S., N. Shephard, and S. Chib (1998). "Stochastic volatility: Likelihood inference and comparison with ARCH models". **Review of Economic Studies** 65, pp. 361–393. DOI: 10.1111/1467-937X.00050 (cit. on pp. 9, 10).

Leydold, J. and W. Hörmann (2014). **GIGrvg: Random variate generator for the GIG distribution**. R package version 0.2. URL: http://CRAN.R-project.org/package=GIGrvg (cit. on p. 39).

References III

McCausland, W. J., S. Miller, and D. Pelletier (2011). "Simulation smoothing for state-space models: A computational efficiency analysis". **Computational Statistics and Data Analysis** 55, pp. 199–212. DOI: 10.1016/j.csda.2010.07.009 (cit. on pp. 11–13).

Omori, Y. et al. (2007). "Stochastic volatility with leverage: Fast and efficient likelihood inference". **Journal of Econometrics** 140, pp. 425–449. DOI: 10.1016/j.jeconom.2006.07.008 (cit. on p. 14).

Rue, H. (2001). "Fast sampling of Gaussian Markov random fields". **Journal of the Royal Statistical Society, Ser. B** 63, pp. 325–338. DOI: 10.1111/1467-9868.00288 (cit. on pp. 11–13).

Yu, Y. and X.-L. Meng (2011). "To center or not to center: that is not the question—An ancillarity-suffiency interweaving strategy (ASIS) for boosting MCMC efficiency". **Journal of Computational and Graphical Statistics** 20, pp. 531–570. DOI: 10.1198/jcgs.2011.203main (cit. on pp. 11–13).