COMPUTATIONAL FINANCE & RISK MANAGEMENT

UNIVERSITY *of* WASHINGTON

Department of Applied Mathematics

The FlexBayes Package

Kjell Konis
Computational Finance & Risk Management
University of Washington

FlexBayes

- FlexBayes is an R package for modeling data in the Bayesian paradigm (Gelman et. al. 2004)
- FlexBayes is a Markov Chain Monte Carlo (MCMC) engine implemented in C++
- Provides functions for fitting the following models

`bhlm`	Hierarchical linear mixed model
`bhpm`	Hierarchical Poisson mixed model
`bhbm`	Hierarchical binomial (logistic) mixed model

Thanks

- TIBCO Software, Inc. for open sourcing FlexBayes
- The coda package (MCMC structures and diagnostics)
- Chapter 7 of Scherer & Martin makes extensive use of FlexBayes

bhlm: Bayesian Hierarchical Linear Model

- bhlm fits a two-stage *hierarchical linear mixed effects* model
- First-level model:

$$y_{ij} \sim \mathcal{N}(\theta_{ij}, \sigma_j^2)$$

[the subscripts refer to observation i in group j]

- Link function:

$$\theta_{ij} = m_{ij}\gamma + x_{ij}\beta_j$$

[γ is fixed effects, β_j is random effects]

- Second-level model:

$$\beta_j = z_j\alpha + u_j$$

where $u_j \sim \mathcal{N}_r(0, \tau^2 V)$

Priors

Parameter	Supported prior distributions
α, γ	normal, t, non-informative (flat)
σ_j^2	inverse chi-squared, uniform shrinkage, non-informative power, DuMouchel
τ^2	inverse chi-squared, uniform shrinkage, non-informative power, DuMouchel
$\tau^2 V$	inverse Wishart

Bayes Alphas and Betas

- x_{ij} empty $\implies$ first-level fixed effects $\implies$ Bayes linear model
- Let $\gamma = (\alpha, \beta)$
- Example: the single-factor market model

$$r_t = \alpha + \beta r_{mt} + \epsilon$$

r_t returns on an individual stock (Microsoft, 2012)
r_{mt} overall market returns (2012)

- Fit the Bayes linear model with non-informative priors

```
> bhlm(rt ~ rmt, data = df)
                    Mean       S.D.
(Intercept) -0.0004702 0.0006032
rmt          1.0750000 0.0792600
SIGMA        0.0098630 0.0004426
```

Bayes Alphas and Betas: non-informative prior

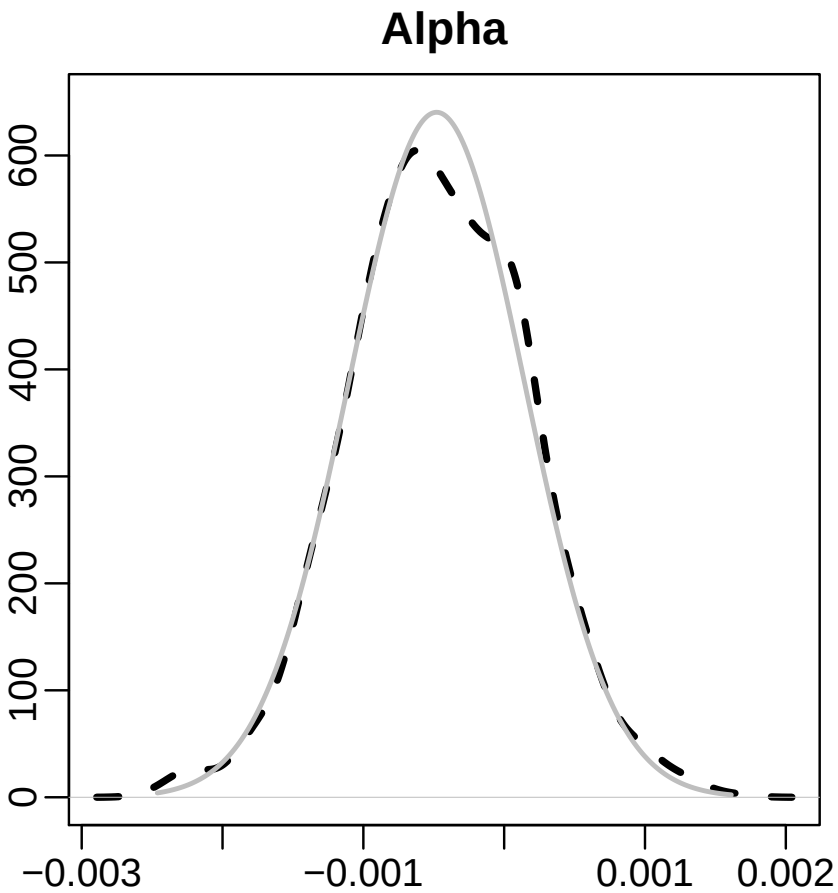

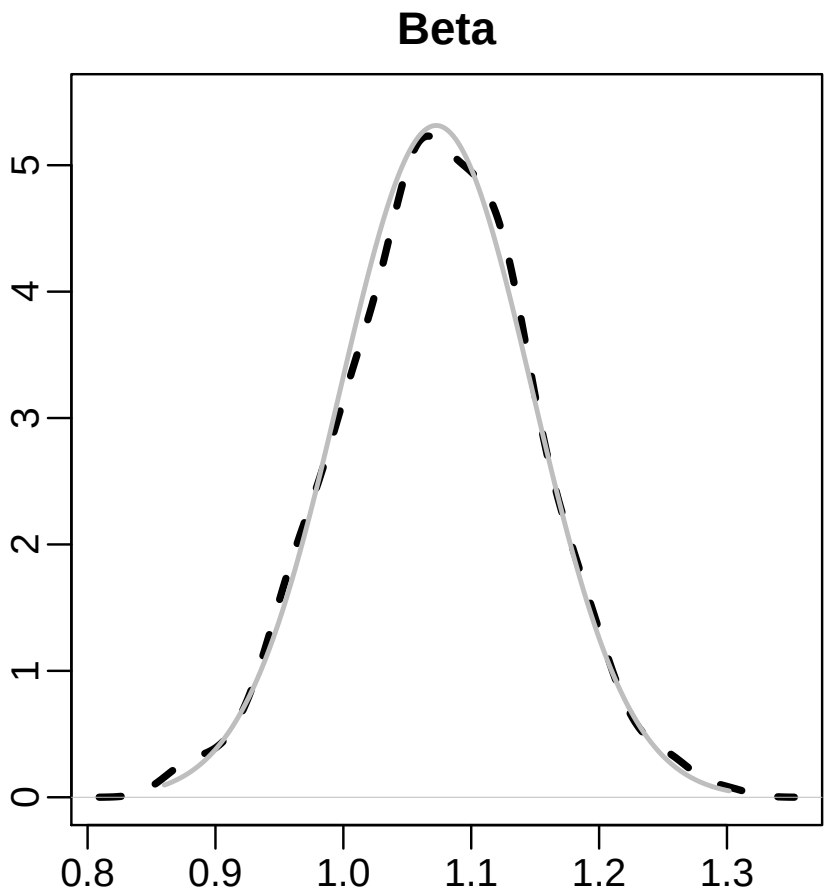

Bayes Alphas and Betas: informative prior

- Informative priors for α and β

$$\begin{bmatrix} \alpha \\ \beta \end{bmatrix} \sim \mathcal{N}_2 \left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0.001^2 & 0 \\ 0 & .1^2 \end{bmatrix} \right)$$

- Fitting the Bayes linear model with informative priors

```
> coef.prior <- fbprior("normal", c(0, 1),
                        diag(c(1e-3^2, 0.1^2)))
> info.prior <- bhlm.prior(fixed = coef.prior)
> bayes.fit <- bhlm(rt ~ rmt, data = df,
                    prior = info.prior)
```

Bayes Alphas and Betas: non-informative prior

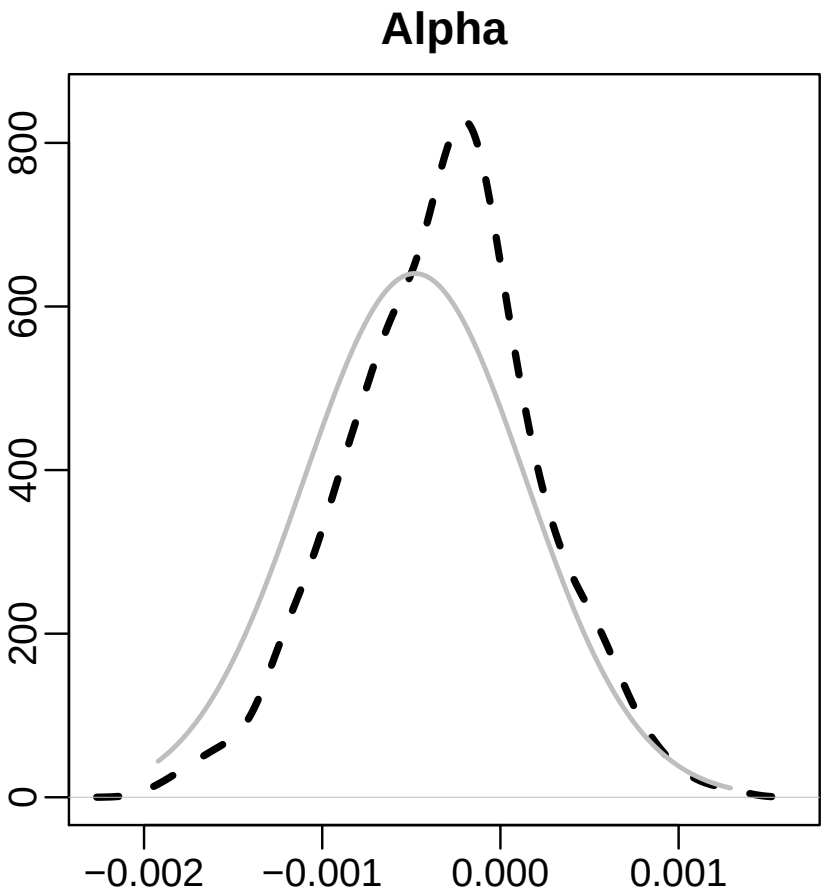

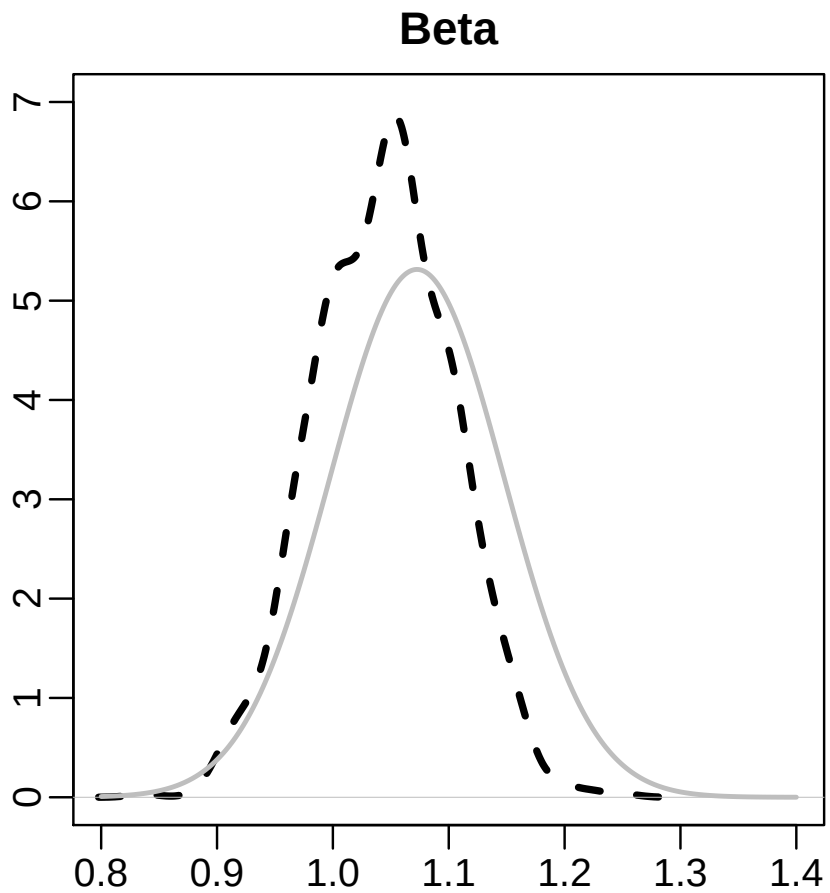

Final Remarks

- The FlexBayes package lives on R-Forge

  http://flexbayes.r-forge.r-project.org

- To install FlexBayes

```
install.packages("FlexBayes",
                 repos="http://R-Forge.R-project.org")
```

References

1. Gelman, A., Carlin, J. B., Stern, H. S. and Rubin, D. B. 2004. *Bayesian data analysis*. Boca Raton, Fla: Chapman & Hall/CRC.
2. Scherer, B. and Martin, D. 2005. *Introduction to Modern Portfolio Optimization with NUOPT and S-PLUS*. New York, NY: Springer.