

MCMCpack: An Evolving R Package for Bayesian Inference

Andrew D. Martin
Washington University
St. Louis, Missouri, USA

Kevin M. Quinn
Harvard University
Cambridge, Massachusetts, USA

<http://mcmcpack.wustl.edu>

20 May 2004



MCMCpack Goals

- Free, open-source, easy-to-use software for Bayesian inference.
- Provide a development environment for easy implementation of non-standard statistical models.
- Provide a distribution mechanism for other researchers with a consistent user interface and documentation.



MCMCpack Design

- Exploit the power of R, making sure the interface is “R-like.”
- Estimation in compiled C++ (using the Scythe Statistical Library), C, or FORTRAN.
- coda mcmc objects to store posterior density samples and simulation parameters.
- coda or boa for posterior density summarization.
- Hidden functions for error checking, passing matrices, etc.

Implemented Models To Date

Linear regression (with Gaussian errors), a general linear panel model, Wakefield's ecological inference model, Quinn's dynamic ecological inference model, Wakefield's hierarchical ecological inference model, a probit model, a logistic regression model, a one-dimensional item response theory model, a K-dimensional item response theory model, a Normal theory factor analysis model, a mixed response factor analysis model, an ordinal item response theory model, a Poisson regression, and an ordered probit model.

An Example: A One Dimensional-IRT Model

```
library(MCMCpack)
data(SupremeCourt)
posterior <- MCMCirtKd(SupremeCourt,
  dimensions=1,
  item.constraints=list("40"=list(2,"+")),
  burnin=1000,
  mcmc=50000,
  thin=5,
  verbose=TRUE,
  B0=1)

print(summary(posterior))
plot(posterior,trace=FALSE)
```

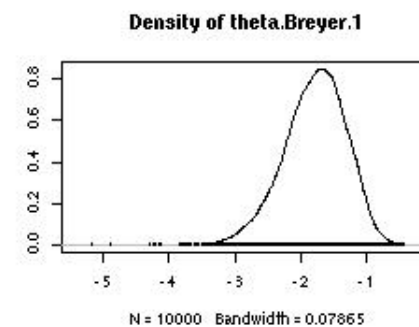
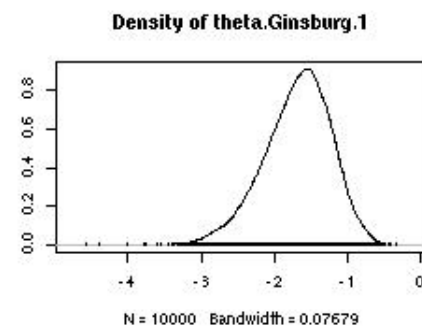
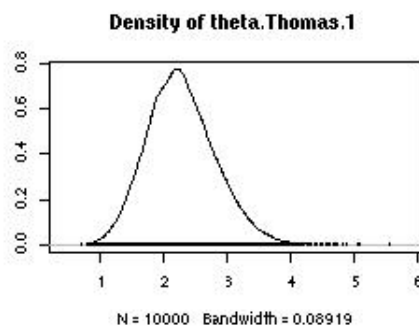
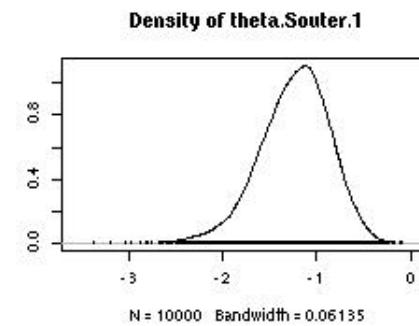
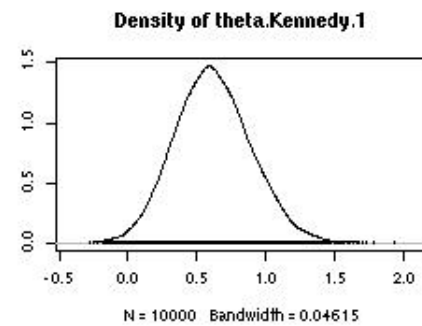
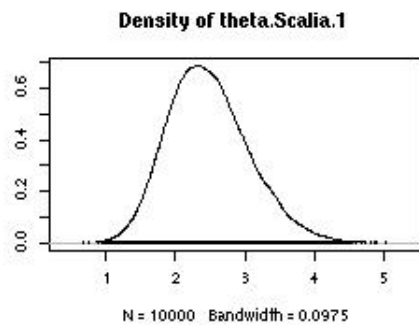
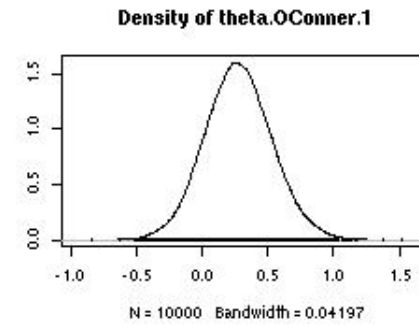
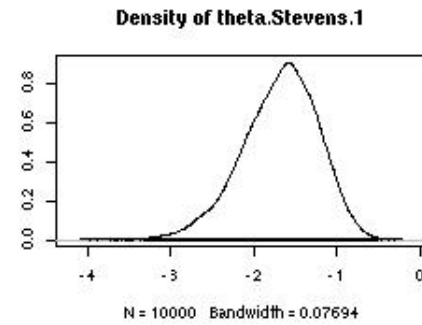
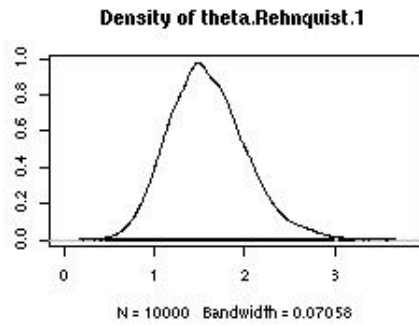
```

adm@tG4 R> source("mcmcpackdemo.R")
##
## Markov chain Monte Carlo Package (MCMCpack)
## Copyright (C) 2003 Andrew D. Martin and Kevin M. Quinn
##
Loading required package: coda
Loading required package: MASS

```

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

	Mean	SD	Naive SE	Time-series SE
theta.Rehnquist.1	1.5945	0.4284	0.004284	0.007341
theta.Stevens.1	-1.6863	0.4605	0.004605	0.008130
theta.OConner.1	0.2752	0.2529	0.002529	0.003560
theta.Scalia.1	2.4801	0.5838	0.005838	0.010581
theta.Kennedy.1	0.6205	0.2776	0.002776	0.004061
theta.Souter.1	-1.2347	0.3702	0.003702	0.006040
theta.Thomas.1	2.2770	0.5488	0.005488	0.009842
theta.Ginsburg.1	-1.6917	0.4632	0.004632	0.008179
theta.Breyer.1	-1.7996	0.4759	0.004759	0.008404



Project Status

- We have been releasing alpha versions for about a year and a half. The package currently contains fourteen models and numerous utility functions.
- Frequent releases and bug fixes.
- Create development environment (complete specification, implement hidden functions and NAMESPACE, document Scythe and MCMCpack API).
- Implement additional models, including “toy” models for classroom instruction.



Things to Think About

- More flexible specification of other priors (not just semi-conjugate priors).
- Flexible estimation engine for other models.
- Real-time visualization of simulation progress.
- GUI?

Acknowledgements

National Science Foundation
Program in Methodology, Measurement, and Statistics
(Grants SES-0350646 and SES-0350613)

Washington University, Department of Political Science and the
Weidenbaum Center on the Economy, Government, and Public Policy
<http://wc.wustl.edu>

Harvard University, Department of Government and the
Center for Basic Research in the Social Sciences
<http://cbrss.harvard.edu>

