Sparse Bayesian Latent Factor Stochastic Volatility Models for Dynamic Covariance Estimation in High-Dimensional Financial Time Series

Gregor Kastner
WU Vienna University of Economics and Business, Austria

Abstract

Dynamic covariance estimation for multivariate time series suffers from the curse of dimensionality; this renders parsimonious approaches essential for conducting reliable statistical inference. We address this issue by modeling the underlying dynamics of a time series vector through a lower dimensional collection of latent factors that allow for time-varying stochastic volatilities.

Furthermore, we apply a Normal-Gamma prior to the elements of the factor loadings matrix. This hierarchical shrinkage prior is a generalization of the Bayesian lasso and effectively pulls the factor loadings of unimportant factors towards zero, thereby increasing sparsity even more.

To guarantee efficiency of the estimation procedure, we employ a fully Bayesian yet computationally feasible approach to obtain draws from the high-dimensional posterior and predictive distributions via Markov chain Monte Carlo (MCMC) samplers. The latent vectors of time-varying volatilities are drawn "all without a loop" (AWOL), and we utilize several variants of an ancillarity-sufficiency interweaving strategy (ASIS) to boost efficiency when sampling the factor loadings as well as the parameters driving the time-varying volatilities.

The effectiveness of the approach is demonstrated through extensive simulation studies. Furthermore, we apply the model to a 20-dimensional exchange rate series and a 300-dimensional vector of stock returns to evaluate predictive performance for financial data.

Keywords: curse of dimensionality, shrinkage, Normal-Gamma prior, Markov chain Monte Carlo (MCMC), ancillarity-sufficiency interweaving strategy (ASIS), all without a loop (AWOL), predictive distribution