Sparse Bayesian Latent Factor Stochastic Volatility Models for High-Dimensional Financial Time Series

Hedibert Freitas Lopes
Insper Institute of Education and Research, Brazil – hedibertFL@insper.edu.br

Dynamic covariance estimation for multivariate time series suffers from the curse of dimensionality, which renders parsimonious estimation methods essential for conducting reliable statistical inference. We address this problem by modeling the underlying volatility dynamics of a time series vector through a lower dimensional collection of latent time-varying stochastic factors. 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 shrinks the factor loadings of unimportant factors towards zero, thereby increasing parsimony even more. Estimation is carried out via Bayesian MCMC methods that make it possible to obtain draws from the high-dimensional posterior distribution. To guarantee efficiency of the samplers, we utilize several variants of an ancillarity-sufficiency interweaving strategy (ASIS) for sampling the factor loadings. We implement the sampler in a compiled programming language, which is interfaced to R for increased usability. Through extensive simulation studies, we demonstrate the effectiveness of the shrinkage prior for sparse loadings matrices. Furthermore, we apply the model to 5000 daily log-returns of 300 stocks listed in the S&P 500 index. This is joint work with Gregor Kastner and Sylvia Fruehwirth-Schnatter from the WU Vienna University of Economics and Business.