

**Abstract:** Many recent papers in macroeconomics have used large Vector Autoregressions (VARs) involving a hundred or more dependent variables. With so many parameters to estimate, Bayesian prior shrinkage is vital in achieving reasonable results. Computational concerns currently limit the range of priors used and render difficult the addition of empirically important features such as stochastic volatility to the large VAR. In this paper, we develop variational Bayes methods for large VARs which overcome the computational hurdle and allow for Bayesian inference in large VARs with a range of hierarchical shrinkage priors and with time-varying volatilities. We demonstrate the computational feasibility and good forecast performance of our methods in an empirical application involving a large quarterly US macroeconomic data set.