Methods for Bayesian Analysis of Time Series. Dissertation Abstract.

Agnieszka Borowska

20.09.2017

Bayesian methods provide a flexible and convenient framework for the analysis of econometric time series mainly due to capturing parameter uncertainty and constituting a natural starting point for model combination. This is of particular importance in the context of risk analysis, where the objective is a precise estimation of rare events and where even a small level of incorrectness might have tremendous and serious consequences. Another advantage of the Bayesian approach is that it highly facilitates dealing with complex data, which otherwise might require imposing of simplifying assumptions. Finally, the Bayesian framework relies to a great extent on simulation based techniques, such as Markov chain Monte Carlo (MCMC) or Sequential Monte Carlo (SMC), which constitutes a convenient starting point for devising non-standard estimation techniques. A big challenge in the context of Bayesian methods is, however, their computational burden. The goal of the research presented in this thesis is to provide accurate and efficient simulation-based methods which can mitigate that problem.

The first chapter of this thesis considers the issue of precise multi-step-ahead risk forecasting. We present an accurate and efficient approach to Bayesian estimation of two financial risk measures, Value at Risk and Expected Shortfall, for a given volatility model. Precise forecasts of the tail of the distribution of returns are obtained not only for the 10-days-ahead horizon required by the Basel Committee but even for long horizons, like one-month or one-year ahead. The latter has recently attracted a considerable attention due to a different character of the short term risk and the long run one. Long horizon forecasts can also be useful e.g. for option pricing. The key insight behind the proposed importance sampling based approach is the construction of the importance densities as mixtures of Student's t distributions sequentially. By oversampling the extremely negative scenarios and punishing them by lower importance weights, a much higher precision in characterising the properties of the left tail is achieved.

In the second chapter we develop a novel estimation and forecasting method for a specific region of the predictive distribution. An important domain of application is accurate prediction of financial risk measures, where the area of interest is the left tail of the predictive posterior density of (log)returns. We specifically concentrate on the Bayesian approach to the issue of model misspecification, which still seems to be an open question under an active, ongoing debate in the Bayesian community. It is known that the posterior concentrates close the points in the support of the prior that minimise the Kullback-Leibler divergence with respect to the true data generating process (DGP). However, from the perspective of accurate tail prediction obtaining estimates being just close to their real values is likely to lead to incorrect risk measures and hence to poor managerial decisions. The proposed method, called the Partially Censored Posterior, originates from the Bayesian approach to parameter estimation and time series forecasting, however it provides a more accurate estimation of the density in the region of interest in case of misspecification. We partition the set of parameters into two subsets: the first, for which we consider the standard marginal posterior, and the second, for which we consider a censored conditional posterior. The censoring means that observations outside the region of interest are censored: for those observations only the probability of being outside the region of interest matters. In the second subset we choose parameters that are expected to benefit from censoring. This approach yields more precise parameter estimation than a fully censored posterior for all parameters, and has more focus on the region of interest than a standard approach. Finally, novel ways of time-varying censoring is developed, beneficial from the tail prediction perspective.

The third chapter presents a method to improve the efficiency of data augmentation algorithms for state-space models. State-space models constitute a broad class of models and provide a flexible framework for studying and forecasting numerous real-life phenomena. This flexibility originates from considering two separate processes: a system process, describing the dynamics of an unobserved (latent) state; and the observation process, which relates the data to the latent component and hence takes into account the observation uncertainty. Data augmentation is a standard approach to perform Bayesian inference, which is able to fit any type of model via imputing all of the unknown states in the complete-data likelihood. However, due to often very high correlation this typically leads to poor mixing of the MCMC algorithm. We propose to circumvent this inefficiency by combining data augmentation with numerical integration in a Bayesian hybrid approach. The underlying idea is to combine the good aspects of both methods but removing the problems that arise. For data augmentation the problem is that of highly correlated unknown states; for numerical integration the problem is that of the curse of dimensionality. To this end, we utilise the structure of the unknown states which can be split into two types: auxiliary variables, which are imputed within the MCMC algorithm using data augmentation; and "integrable" states, which are numerically integrated out within the likelihood expression. The idea is to specify the unknown states in such a way that the algorithm is efficient. The proposed technique can be applied to different types of problems including estimation of the stochastic volatility for financial data or abundance estimation for ecological time series.