Summary

This thesis concerns the state space approach to time series analysis as a flexible framework for modelling time-changing parameters in economics and finance. We focus on the methodological issues related to the simulation-based methods which are necessary for analysing these models outside the simplest settings. We also develop new econometric models that can increase our understanding of the time-varying nature of economic risks and the instability of financial parameters.

In Chapter 2 we introduce a new efficient importance sampler for nonlinear non-Gaussian state space models. We propose a general and efficient likelihood evaluation method for this class of models via the combination of numerical and Monte Carlo integration methods. Our methodology explores the idea that only a small part of the likelihood evaluation problem requires simulation. We refer to our new method as numerically accelerated importance sampling (NAIS). We show that the NAIS method is computationally and numerically efficient, facilitates parameter estimation for models with high-dimensional state vectors, and overcomes a bias-variance trade-off encountered by other sampling methods. An elaborate simulation study for stochastic volatility, stochastic conditional duration and stochastic copula models as well as an empirical application for U.S. stock returns reveal large efficiency gains for a range of models used in financial econometrics.

In Chapter 3 we develop a systematic framework for linking a general class of discrete time stochastic volatility (SV) models to realised measures of volatility such as the two time scales estimator of Zhang, Mykland, and Aït-Sahalia (2005), the realised kernel of Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008) and the pre-averaging based realised variance estimator of Jacod, Li, Mykland, Podolskij, and Vetter (2009). Our analysis considers a fully specified time series model for both the returns and the realised measures. We assume a linear state space representation for the log realised measures, which are noisy and biased estimates of the log integrated variance, at least due to Jensen's inequality. We incorporate filtering methods for the estimation of the latent log volatility process.

We contribute to the literature by recognising that the dependence between daily returns and measurement errors affects the estimation of the model. We develop a two-step estimation method for the parameters in our specification which overcomes this

problem. This method is computationally straightforward even when the stochastic volatility model contains non-Gaussian return innovations and leverage effects. We perform a detailed empirical study of the realised SV model using data for nine Dow Jones index stocks in the period between 2001 and 2010. We find that measurement errors account for between 24% and 53% of the variance of daily innovations in the log realised kernel and pre-averaging based realised variance series. We show that time series filtering leads to important reductions in the variance of the log volatility estimates. We also find that forecasts from our model outperforms those from a set of recently developed alternatives.

In Chapter 4 we study whether and when parameter-driven time-varying parameter models lead to forecasting gains over observation-driven models. We consider dynamic count, intensity, duration, volatility and copula models, including specifications which have not been studied earlier in the literature. In an extensive Monte Carlo study, we find that observation-driven generalised autoregressive score (GAS) models have similar predictive accuracy to correctly specified parameter-driven models. In most cases, differences in mean squared errors are smaller than 1%, so that model confidence sets (Hansen, Lunde, and Nason 2011) have low power when comparing these two alternatives. We also find that GAS models outperform many familiar observation-driven models in terms of forecasting accuracy. The results point to a class of observation-driven models with comparable forecasting ability to parameter-driven models, but lower computational complexity.

In chapter 5 we return to the methodological questions which first appeared in Chapter 2. We propose a new likelihood evaluation method based on importance sampling for nonlinear non-Gaussian state space models with multiple time-varying parameters in the observation density. Our sampler consists of an efficient approximating linear state space model for which we estimate optimal parameters by solving series of low dimensional integrals using a quasi-Monte Carlo method. We develop new results that substantially reduce the computational complexity of the algorithm as a function of the number of time-varying parameters in the state space model. We provide several Monte Carlo and empirical illustrations of our method in challenging settings. We consider stochastic volatility models with leverage effects and non-Gaussian return innovations and a multivariate stochastic volatility specification. Our results reveal up to 95% gains in the time normalised variance of the likelihood estimates over a standard efficient importance sampling (EIS) approach. We also show that the method leads to small Monte Carlo errors for parameter estimation in practice, even in large sample sizes.

References

Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008). Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica 76*, 1481–1536.

Hansen, P. R., A. Lunde, and J. M. Nason (2011). The Model Confidence Set. *Econometrica 79*(2), 453–497.

Jacod, J., Y. Li, P. A. Mykland, M. Podolskij, and M. Vetter (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and their Applications 119*(7), 2249–2276.

Zhang, L., P. Mykland, and Y. Aït-Sahalia (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association 100*, 1394–1411.