Abstract
The first chapter focuses on evaluation of time-series forecasts. It is a common practice to split a time series into in-sample and pseudo out-of-sample segments and estimate the out-of-sample loss for a given statistical model by evaluating forecasting performance over the pseudo out-of-sample segment. I propose an alternative estimator of the out-of-sample loss, which, contrary to conventional wisdom, utilizes criteria measured both in- and out-of-sample via a carefully constructed system of affine weights. I prove that, provided the time series is stationary, the proposed estimator is the best linear unbiased estimator of the out-of-sample loss, and outperforms the conventional estimator in terms of sampling variability. Application of the optimal estimator to Diebold-Mariano type tests of predictive ability leads to a substantial power gain without increasing finite sample size distortions. An extensive evaluation on real world time series from the M4 forecasting competition confirms the superiority of the proposed estimator, and also demonstrates substantial robustness to violations of the underlying assumption of stationarity.

In the second chapter we perform an extensive investigation of different specifications of the BEKK-type multivariate volatility models for a moderate number of assets, focusing on how the degree of parametrization affects forecasting performance. Because the unrestricted specification may be too generously parameterized, often one imposes restrictions on coefficient matrices constraining them to have a diagonal or even scalar structure. We frame all three model variations (full, diagonal, scalar) as special cases of a ridge-type regularized estimator, where the off-diagonal elements are shrunk towards zero and the diagonal elements are shrunk towards homogeneity. Our forecasting experiments with the BEKK-type Conditional Autoregressive Wishart model for realized volatility confirm the superiority of the more parsimonious scalar and diagonal model variations. Regularization of the diagonal and off-diagonal parameters does not regularly lead to tangible performance improvements irrespective of how precise the tuning of regularization intensity is. Additionally, our results highlight the crucial importance of frequent model re-estimation in improving the forecast precision, and, perhaps paradoxically, a slight advantage of shorter estimation windows compared to longer windows.

In the third chapter I propose a novel meta-learning model that utilizes hypernetworks to design a parametric model tailored to a specific family of forecasting tasks. The model's training can be directly performed with backpropagation, eliminating the need for reliance on higher-order derivatives, and is equivalent to a simultaneous search over the space of parametric functions and their optimal parameter values. This, in essence, provides a data-driven alternative to manually designing a parametric model for a group of similar prediction tasks, an endeavor that typically requires considerable statistical expertise and domain knowledge. I demonstrate the capabilities of the proposed meta-learning model on two applications. When applied to the sinusoidal regression task, the proposed model outperforms state-of-the-art meta-learning approaches and is capable of almost perfectly recovering the underlying parametric model. As a second application, I use the model to make quintile predictions for asset returns in the M6 Financial Forecasting Competition. The model attained an RPS of 0.15689, securing the 4th place in the forecasting challenge and ultimately the 1st place in the overall duathlon ranking.