Dynamic Covariance Modelling Using Generalised Wishart Processes

When constructing a financial portfolio it is important to have an idea of how the different assets are correlated in order to make a fair assessment of the risk involved. Unfortunately, finding good covariance estimates has proven a difficult task across many fields. In this project, we set up a framework to obtain a distribution of covariances, rather than the single estimate typically produced by traditional techniques, with filters arguably being the most common. The resulting estimates react well to underlying changes, and carry a selection of desirable properties over the more *engineery* methods. We run the model on both synthetic and real data and, although the test scenarios are limited, the approach seems promising. An existing variant of the framework is extended by implementing a more more efficient algorithm, making use of dynamics found in quantum mechanics.

With a proper understanding of the correlation structure, the portfolio can be designed to mitigate risk – this is known as hedging. It has however proven difficult to estimate the correlation between different assets, and to make matters worse it is typically not constant over time as market conditions change. To account for this time dependence we need to build a *dynamic* modelling framework, which is exactly what this project aims to do. A very common method uses *filters* which are functions mapping the observed data onto some estimate of the covariance, oftentimes designed to favour more recent observations. Instead of following this ad-hoc design, we take a more probabilistically rigorous approach, and set out to find a *distribution* of the covariance given the observed data, rather than just a single estimate.

The idea of assuming some prior model for the covariance and then updating it depending on the data is known as *Bayesian inference*, and is one of two dominant statistical paradigms. We perform the inference by adapting a *Markov chain Monte Carlo* method, which is very common when handling difficult distributions such as the one we are looking for. We utilise two staple methods known as *Metropolis-Hastings* and *Gibbs sampling*, but also deploy two more intricate methods used specifically for our purpose, called *Elliptical slice sampling* and *Hamiltonian Monte Carlo*. The latter forms an interesting bridge between quantum mechanics and probability by evolving a system of statistical entities under Hamiltonian dynamics.

In reality we cannot observe the true covariance, so to validate the method we first test the framework on synthetic data generated from a covariance which we know, with varying complexity. With access to the true covariance, we set up a few measurements of quantifying how far away the sample covariances lie, and compare this to a filter-based benchmark model. We then deploy the model on financial time series and evaluate the results using a selection of common portfolio metrics. While the results are somewhat mixed, performance is broadly speaking in parity with the benchmarks. It would be natural to perform some more extensive testing to further warrant the framework, after which there are many ways to extend and improve the model.

A thesis by Fredrik Nilsson, spring 2023. Department of Mathematical Statistics, Lund Institute of Technology. Written in collaboration with Lynx Asset Management AB, Stockholm.