Statistical methods for high-dimensional volatility

Financial volatility has been a central topic in statistical research for several decades. Classical high-frequency methods—most notably realized volatility—provide consistent estimators of price process volatility, and their asymptotic properties are now well understood thanks to seminal contributions by Jacod and co-authors.

However, modern financial applications increasingly involve large portfolios, where both the dimension of the data and the size of the portfolio can grow together. This raises new challenges for volatility estimation in high-dimensional settings.

In this talk, we present recent advances in statistical inference for high-dimensional volatility. We focus on volatility estimation under low-rank structures and investigate the asymptotic behaviour of the spectral distribution of the integrated volatility matrix. Our methodology combines a range of probabilistic tools, including stable central limit theorems, concentration inequalities, and techniques from random matrix theory.