

Shrinking the Covariance Matrix: A Portfolio Perspective

Rodolphe Hadrien Vanderveken

UCLouvain
Voie du Roman Pays 34/L1.03.01 1348 Louvain-la-Neuve
Belgium

Joint work with: Nathan Lassance, Frédéric Vrins

Estimating the covariance matrix is a central problem in portfolio selection. The foundational shrinkage methodologies developed by Ledoit and Wolf (2004, 2017) suffer from two drawbacks: they are not designed to optimize out-of-sample portfolio performance and do not account for estimation errors in the means. In this paper, we propose a novel shrinkage covariance matrix estimator that addresses these two drawbacks. Specifically, we calibrate the shrinkage intensities in linear and nonlinear shrinkage estimators so that they maximize the expected out-of-sample portfolio performance. We find that this alternative calibration results in higher shrinkage intensities relative to the traditional approach and delivers a superior out-of-sample portfolio performance. Overall, our methodology is a one-step approach that estimates the covariance matrix and the optimal portfolio at the same time, which delivers large economic gains relative to the conventional two-step scheme.