

While the use of statistical methods to identify financial risk factors is a long-standing practice, the use of convex optimization for this purpose is a recent innovation. Specifically, a combination of convex programs developed by Chandrasekaran, Parillo and Witsky (2012) and Saunderson et al. (2012) can be used to extract financial risk factors from a sample return covariance matrix. I will outline an application of this extraction to financial risk forecasting and develop finance-oriented metrics to assess accuracy. Finally, I provide an example that highlights the difference between this approach and principal component analysis, which is the academic standard for risk factor identification. In a companion talk, Alex Shkolnik will discuss the convergence properties of the convex programs, and he will analyze the accuracy of the algorithm on simulated data.