Editorial

Modern finance has become increasingly technical, requiring the use of sophisticated mathematical tools in both research and practice. Many find the roots of this trend in the portfolio selection models and methods described by Markowitz in the 1950s and the option pricing formulas developed by Black, Scholes and Merton in the late 1960s. For quantitative investment management, Markowitz' mean-variance optimisation (MVO) models have had an enduring effect and continue to frame the research and discussion on portfolio management to this day.

Optimisation is a mature branch of applied mathematics where the main objective is to allocate scarce resources to alternative activities in a way that maximises the combined benefits of these activities. Decades of intensive research led to the discovery of fast and reliable algorithms and software for the solution of many important classes of optimisation problems. As a result, optimisation can be and is used as an effective management tool in various industries, including investment management. The paper by Khodadadi *et al.* in this special issue studies the role of optimisation in an active investment management process, addressing many issues from building an investment management platform to handling transaction costs. This paper also investigates challenges in the use of optimisation as an effective investment management tool touching on the issues of constraints, multi-period models, parameter estimation risk, among others. Mulvey, Simsek and Zhang's

contribution in this issue demonstrates the benefits of multi-period investment models in pension fund management using stochastic optimisation techniques. Drawing on examples from the US telecommunication sector and an underfunded hypothetical car company, they develop overlay strategies as well as strategies that change allocations based on certain triggers, resulting in significant improvements in the fund's value in simulated tests.

Despite their effective use in investment management, portfolio optimisation methods never lacked critics or sceptics. It can be argued that part of the reason for this scepticism is our lack of understanding of the complex relationship between the inputs and outputs of an optimisation problem. In often-cited examples, researchers and portfolio managers reported that optimal portfolios change drastically and inexplicably in response to small changes in the problem inputs such as alphas. Such examples fuelled the distrust for optimisers and served as an impetus to the pioneering work by Black and Litterman (1992) and, more recently, the emergence of robust optimisation techniques. Robust optimisation explicitly models the potential errors and uncertainties in the inputs of the optimisation problem to efficiently generate outputs that are less sensitive to these input errors. Papers by Ceria and Stubbs as well as by Schöttle and Werner in this issue explore the use of robust optimisation methodologies in portfolio construction. In extensive computational

tests, Ceria and Stubbs show that portfolios generated using the robust optimisation methodology typically outperform traditional mean-variance optimised portfolios and are more stable. Similarly, Schöttle and Werner compare robust optimisation methods with simulation-based re-sampling techniques and report that both approaches improve out-of-sample performance of portfolios. They also provide theoretical and empirical evidence for the stability of robust optimal portfolios when inputs are subject to perturbations.

One of the other important developments in theoretical and applied optimisation during the last decade was the explosion of research in conic and other convex optimisation problems. ‘Fancy’ problem classes such as second-order cone optimisation and semi-definite optimisation were intensely studied, resulting in efficient algorithms and reliable software. These problem classes are now being introduced into financial models. For example, the robust optimisation models studied by Ceria and

Stubbs as well as by Schöttle and Werner result in second-order cone optimisation problems and were solved using these advanced techniques. Ndiaye *et al.* provide an interesting application of conic optimisation models in the final paper of this special issue. They address the problem of building a global risk model from regional risk models and cross-region correlations in a consistent way and show that this problem can be formulated as a semi-definite optimisation problem and solved efficiently using advanced optimisation algorithms.

Optimisers are sophisticated tools, and a good deal of practical experience is often required to ensure that the investment manager is maximising their value. We believe the five papers in this special issue provide excellent guidance in this direction and serve to advance our understanding of the application of optimisation methods to investment management. Enjoy!

Reha H. Tütüncü and Peter J. Zangari
Guest Editors