

Portfolio Optimization with R/Rmetrics

For the past several years, Diethelm Würtz and his group at the Econophysics Group at the Institute of Theoretical Physics of ETH Zurich has been implementing tools in [R](#) for financial and econometric analysis. The [Rmetrics project](#) has now swelled to more than 30 packages for R, all available for [free download](#).

Despite the fact Rmetrics is used in an industry which uses it to generate untold wealth (most of the time, at least), the [Rmetrics Association](#) which produces and supports Rmetrics is run as a non-profit enterprise. An affiliated organization, Finance Online, does generate financial support for Rmetrics through consulting to financial firms, but otherwise funding comes through donations.

Now, there's new way to contribute to the project by purchasing an E-book describing how to use Rmetrics. A series of E-books is planned, but the first is available now: [Portfolio Analysis with R/Rmetrics](#). At around \$85 (US) it's competitively priced with printed volumes that cover similar topics, and comes in a convenient PDF format. There's no copy-protection associated with the PDF (other than your name printed on each page to discourage unpaid distribution), and your purchase price includes all updates to the volume for a year. The PDF itself is very well produced, thanks to the [talents](#) of producer and co-author Andrew Ellis. The examples and output are all automatically generated from R and integrated with the written content to ensure accuracy, and the text is extensively bookmarked and hyperlinked making it easy to skip from section to section on-screen.

At over 450 pages it's a comprehensive study of all aspects of portfolio optimization with Rmetrics. If you're new to the domain (but have a good grounding in statistics and analysis), the theory sections provide a welcome and concise overview to the methods implemented. It does assume some familiarity with R, but all examples all start from first principles and include clear and well-commented code. Data sets are provided for the examples, and instructions are also provided on how to download financial data from public sources like Yahoo, the Swiss Exchange, and the Federal Reserve Bank in St Louis if you want to work with up-to-the-minute data sets.

There's a detailed review of the contents of the book after the break, but if you're doing any kind of financial analysis with R and haven't yet looked at Rmetrics, this book is a great place to start.

Rmetrics Association: [Portfolio Optimization with R/Rmetrics](#)

Contents of *Portfolio Analysis with R/Rmetrics*

The book is divided into six major parts, each with its own chapters.

Part I is concerned with creating and managing financial data. It covers how to deal with dates, times, and time series data in R. Examples demonstrate how to combine, align, merge and aggregate assets in multiple time series, and how to select subsets or apply rolling calculations. Computations on time series are next: how to calculate returns, drawdowns and durations, and how to calculate statistics including univariate measures, the sample covariance matrix, and higher-order moments. The part concludes with a detailed section on robust mean and covariance estimates.

Part II focuses on exploratory analysis of financial data. It begins with how to plot time series (especially multiple time series on the same time axis). Next it covers how to graphically explore distributions: box plots, histograms, Q-Q plots, and kernel density estimators. For presentation graphics, there's a chapter on customizing the appearance of plots (including a detailed section on selecting colors). The final sections move more into analytical techniques: simulating assets and testing distributional assumptions. There are also chapters on selecting similar or dissimilar assets with clustering techniques and eigenvalue analysis, and looking at pairwise dependence of assets.

Part III introduces the data structures in Rmetrics for representing portfolios, portfolio data, portfolio constraints (including long-only, unlimited short selling, box constraints, group constraints, risk budgets, and nonlinear and user-defined constraints), and the functions for operating on these portfolio objects.

Part IV focuses on mean-variance portfolio theory and how to specify such portfolios and their constraints. It then goes on to show how to optimize mean-variance portfolios in various ways, including how to optimize the Sharpe Ratio and how to find and display the efficient frontier. This section includes a detailed case study of optimizing a portfolio comprising the 30 elements of the Dow Jones Index. There's also a section on robust portfolio optimization.

Part V moves onto mean-CVaR (conditional value-at-risk) portfolios: how to specify them, and how optimize (including calculating the maximum risk/return portfolio and various kinds of portfolio frontier).

Part VI delves into the practical application of backtesting to portfolio strategies, including how to specify historical windows and your own backtesting function to apply to them, and how to display the backtesting results graphically. The section concludes with two detailed case studies: backtesting the sector rotation of a portfolio of Swiss equities; and backtesting GCC Index rotation.

The book is rounded out with some useful appendices: a review of the various R packages used in the book material; descriptions of the data sets used, and a detailed bibliography.

David Smith, August 2th, 2009,
<http://blog.revolution-computing.com/2009/08/portfolio-analysis-with-rrmetrics.html>