

R documentation

of 'man/PerformanceAnalytics-package.Rd'

March 5, 2007

PerformanceAnalytics-package

Econometric tools for performance and risk analysis.

Description

Library of econometric functions for performance and risk analysis of financial portfolios. This library aims to aid practitioners and researchers in using the latest research in analysis of both normal and non-normal return streams.

We created this library to include functionality that has been appearing in the academic literature on performance analysis and risk over the past several years, but had no functional equivalent in R. In doing so, we also found it valuable to have wrapper functions for functionality easily replicated in R, so that we could access that functionality using a function with defaults and naming consistent with common usage in the finance literature. The following sections cover Performance Analysis, Risk Analysis (with a separate treatment of VaR), Summary Tables of related statistics, Charts and Graphs, a variety of Wrappers and Utility functions, and some thoughts on work yet to be done.

Details

Package: PerformanceAnalytics
Type: Package
Version: 0.9.3
Date: 2007-03-05
License: GPL

In general, this library is most tested on return (rather than price) data on a monthly scale. Many functions will work with daily or irregular return data as well. See function [CalculateReturns](#) for calculating returns from prices, and be aware that the [aggregate](#) function has methods for `tseries` and `zoo` timeseries data classes to rationally coerce irregular data into regular data of the correct periodicity.

In this summary, we attempt to provide an overview of the capabilities provided by `PerformanceAnalytics` and pointers to other literature and resources in R needed for performance and risk analysis. We hope that this summary and the accompanying package and documentation partially fill a hole in the tools available to a financial engineer or analyst.

Performance Analysis

Performance measurement is the first step down a longer road. When discussing performance at an abstract level, terms such as “return” and “risk” are deliberately vague because they help to disguise some (usually unnecessary) complexity. Scratch the surface to answer more complicated questions, and that complexity is necessarily uncovered again. Indeed, the literature around the subject of performance analysis seems to have exploded with the popularity of alternative assets such as hedge funds, managed futures, commodities, and structured products. Simpler tools that seemed appropriate in a relative investment world seem inappropriate for an absolute return world. Risk measurement, which is nearly inseparable from performance assessment, has been on a long march to become multi-dimensional and multi-moment while trying to answer the question: “How much could I lose?” Portfolio construction and risk budgeting are two sides of the same coin: “How do I maximize my expected gain and avoid going broke?” But before we can approach those questions we first have to ask: “Is this something I might want in my portfolio?”

With the increasing availability of complicated alternative investment strategies to both retail and institutional investors, and the broad availability of financial data, an engaging debate about performance analysis and evaluation is as important as ever. There won't be one *right* answer delivered in these metrics and charts. What there will be is an accretion of evidence, organized to *assist* a decision maker in answering a specific question that is pertinent to the decision at hand. Using such tools to uncover information and ask better questions will, in turn, create a more informed investor.

Performance measurement starts with returns. Proprietary traders will object, complaining that “You can't eat returns,” and will prefer to look at numbers with currency signs. To some extent, they have a point - the normalization inherent in calculating returns can be deceiving. Most of the recent work in financial econometrics, however, is focused on returns rather than prices. This “price per unit of investment” standardization is important for two reasons - first, it helps the decision maker to compare opportunities, and second, it has some useful statistical qualities. As a result, the `PerformanceAnalytics` library focuses on returns. See `CalculateReturns` for converting net asset values or prices into returns, either discrete or continuous. See `portfolioReturn` in package `fPortfolio` for calculating returns from a portfolio of instruments with a combination of historical returns and a weighting vector.

Returns and risk may be annualized as a way to simplify comparison over longer time periods. Although it requires a bit of estimating, such aggregation is popular because it offers a reference point for easy comparison. Examples are in `Return.annualized`, `StdDev.annualized`, and `SharpeRatio.annualized`.

Basic measures of performance tend to treat returns as independent observations. In this case, the entirety of R's base is applicable to such analysis. Some basic statistics we have collected in `table.MonthlyReturns` include:

<code>mean</code>	arithmetic mean
<code>mean.geometric</code>	geometric mean
<code>mean.stderr</code>	standard error of the mean (S.E. mean)
<code>mean.LCL</code>	lower confidence level (LCL) of the mean
<code>mean.UCL</code>	upper confidence level (UCL) of the mean
<code>quantile</code>	for calculating various quantiles of the distribution
<code>min</code>	minimum return
<code>max</code>	maximum return
<code>range</code>	range of returns
<code>length(R)</code>	number of observations
<code>sum(is.na(R))</code>	number of NA's

Modern Portfolio Theory (MPT)

Modern Portfolio Theory is the collection of tools and techniques by which a risk-averse investor

may construct an optimal portfolio. It was pioneered by Markowitz's ground breaking 1952 paper *Portfolio Selection*. It also encompasses CAPM, below, the efficient market hypothesis, and all forms of quantitative portfolio construction and optimization.

The *Capital Asset Pricing Model (CAPM)*, initially developed in *William Sharpe in 1964*, provides a justification for passive or index investing by positing that assets that are not on the efficient frontier will either rise or lower in price until they are on the efficient frontier of the market portfolio. The `CAPM.RiskPremium` on an investment is the measure of how much the asset's performance differs from the risk free rate. Negative Risk Premium generally indicates that the investment is a bad investment, and the money should be allocated to the risk free asset or to a different asset with a higher risk premium. `CAPM.alpha` is the degree to which the assets returns are not due to the return that could be captured from the market. Conversely, `CAPM.beta` describes the portions of the returns of the asset that could be directly attributed to the returns of a passive investment in the benchmark asset. The Capital Market Line `CAPM.CML` relates the excess expected return on an efficient market portfolio to its risk (represented in CAPM by `StdDev`). The slope of the CML, `CAPM.CML.slope`, is the Sharpe Ratio for the market portfolio. The Security Market Line is constructed by calculating the line of `CAPM.RiskPremium` over `CAPM.beta`. For the benchmark asset this will be 1 over the risk premium of the benchmark asset. The slope of the SML, primarily for plotting purposes, is given by `CAPM.SML.slope`. CAPM is a market equilibrium model or a general equilibrium theory of the relation of prices to risk, but it is usually applied to partial equilibrium portfolios, which can create (sometimes serious) problems in valuation.

The performance premium provided by an investment over a passive strategy (the benchmark) is provided by `ActivePremium`, which is the investment's annualized return minus the benchmark's annualized return. A closely related measure is the `TrackingError`, which measures the unexplained portion of the investment's performance relative to a benchmark. The `InformationRatio` of an Investment in a MPT or CAPM framework is the Active Premium divided by the Tracking Error. Information Ratio may be used to rank investments in a relative fashion.

We have also included a function to compute the `KellyRatio`. The Kelly criterion, properly applied, will maximize log-utility of returns and avoid risk of ruin. Utilizing the Kelly Criterion to determine leverage or investment size on a strategy is guaranteed to avoid risk of ruin and eventually (over a long enough time horizon) create wealth, but it may be incredibly volatile, with a N% chance of being down N% at some point in time. Even when an investor or analyst does not intend to utilize the Kelly criterion as part of the investment sizing strategy, it can be used as a stack-ranking method like `InformationRatio` to describe the "edge" an investment would have over a random strategy or distribution.

`InformationRatio`, `KellyRatio`, `SharpeRatio`, `SortinoRatio`, Spearman Rank Correlation provided by `rcorr` and other methods discussed here are all methods of rank-ordering relative performance. *Alexander and Dimitriu(2004)* in "*The Art of Investing in Hedge Funds*" show that relative rankings across multiple pricing methodologies may be positively correlated with each other and with expected returns. This is quite an important finding because it shows that multiple methods of predicting returns and risk which have underlying measures and factors that are not directly correlated to another measure or factor will still produce widely similar quantile rankings, so that the "buckets" of target instruments will have significant overlap. This observation specifically supports the point made early in this document regarding "accretion of the evidence" for a positive or negative investment decision.

It is often valuable when evaluating an investment to know whether the instrument that you are examining follows a normal distribution. One of the first methods to determine how close the asset is to a normal or log-normal distribution is to visually look at your data. Both `chart.QQPlot` and `chart.Histogram` will quickly give you a feel for whether or not you are looking at a normally distributed return history. Differences between `var` and `SemiVariance` will help you identify `skewness` in the returns. Skewness measures the degree of asymmetry in the return distribution. Positive Skewness indicates that more of the returns are positive, negative skewness indicates that

more of the returns are negative. An investor should in most cases prefer a positively skewed asset to a similar (style, industry, region) asset that has a negative skewness. Kurtosis measures the concentration of the returns in any given part of the distribution (as you should see visually in a histogram). The `kurtosis` function will by default return what is referred to as “excess kurtosis”, where 0 is a normal distribution, other methods of calculating kurtosis than `method="excess"` will set the normal distribution at a value of 3. In general a rational investor should prefer an asset with a low to negative excess kurtosis, as this will indicate more predictable returns. If you find yourself needing to analyze the distribution of complex or non-smooth asset distributions, the `nortest` package has several advanced statistical tests for analyzing the normality of a distribution.

One question that an investor or researcher is often required to answer revolves around how to diversify a portfolio. In a mean-variance world, an asset is a diversifier if it lowers the variance (and hopefully also raises the mean return) of the portfolio. In a world with more nuanced risk metrics, you may also test the target portfolio to see if it has lower measures of risk than the current portfolio. If you want to do more analysis, we have provided functions to calculate the higher moments and co-moments of the distribution. We have already discussed `skewness` and `kurtosis` provided by package `fBasics`. Increasing skewness and decreasing kurtosis of the target portfolio probably marks a good diversifier. You can also work on minimizing the co-moments and systematic co-moments of the distribution to improve the resilience of the portfolio to multiple different market conditions in much the way that minimizing covariance and systematic beta will mark a well diversified portfolio in a mean-variance framework. See `CoSkewness` and `CoKurtosis` for the co-moments and `SystematicSkewness` and `SystematicKurtosis` for the systematic standardized co-moments. Other diversification potential in a portfolio is not strictly econometric, and follow the classic asset delineations of asset type, size, industry, growth/value, currency, or geographic region, all of which are well-covered elsewhere.

We have covered multiple methods of performance *analysis* in this summary section, but it is also important to note that we have *not* covered performance *attribution*. We have not examined tools for determining the sources and causes for the returns, or the underlying causes for the risks of an instrument (more on Risk Measurement and prediction below). There is a significant amount of academic literature on identifying these sources of risk or returns. Much of it falls into the field of “factor analysis” where “risk factors” are used to retrospectively explain sources of risk, and through regression and other analytical methods *predict* future period returns and risk based on factor drivers. These are well covered in chapters on factor analysis in *Zivot and Wang(2006)* and also in the R functions `factanal` for basic factor analysis and `princomp` for Principal Component Analysis. The authors feel that financial engineers and analysts would benefit from some wrapping of this functionality focused on finance, but the capabilities already available from the base functions are quite powerful.

Summary Tabular Data

Summary statistics are then the necessary aggregation and reduction of these (potentially thousands) of periodic return numbers. Usually these statistics are most palatable when organized into a table of related statistics, assembled for a particular purpose. A common offering of past returns organized by month and cumulated by calendar year is usually presented as a table, such as in `table>Returns`. Adding benchmarks or peers alongside the annualized data is helpful for comparing returns in calendar years.

When we started this project, we debated whether or not such tables would be broadly useful or not. No reader is likely to think that we captured the precise statistics to help their decision. We merely offer these as a starting point for creating your own. Add, subtract, do whatever seems useful to you. If you think that your work may be useful to others, please consider sharing it so that we may include it in a future version of this library.

Other tables for comparison of related groupings of statistics discussed elsewhere in this summary:

<code>table.MonthlyReturns</code>	provides a statistical summary of the monthly returns themselves
<code>table.AnnualizedReturns</code>	calculates aggregate returns and risk (usually used in graphics, see below)
<code>table.CAPM</code>	a collection of statistics offered by practitioners of Modern Portfolio Theory
<code>table.HigherMoments</code>	metrics for evaluating diversification potential by comparing higher moments
<code>table.DownsideRisk</code>	statistics for those wanting to compare the extent of observed losses
<code>table.Correlation</code>	a table arranged for comparing correlations
<code>table.RollingPeriods</code>	calculates statistics over varying periods of time

Charts and Graphs

Graphs and charts can also help to organize the information visually. The cumulative returns or wealth index is usually the first thing displayed, even though neither conveys much information. See `chart.CumReturns`. Individual period returns may be helpful for identifying problematic periods, such as in `chart.Bar`. Risk measures can be helpful when overlaid on the period returns, to display the bounds at which losses may be expected. See `chart.BarVaR` and the following section on Risk Analysis. More information can be conveyed when such charts are displayed together, as in `charts.PerformanceSummary`, which combines the performance data with detail on downside risk (see `chart.Drawdown`).

Two-dimensional charts can also be useful while remaining easy to understand. `chart.Scatter` is a utility scatter chart with some additional attributes that are used in `chart.RiskReturnScatter`. Overlaying Sharpe ratio lines or boxplots helps to add information about relative performance along those dimensions. The relative performance through time of two assets can be plotted with `chart.RelativePerformance`. This plot displays the ratio of the cumulative performance at each point in time and makes periods of under- or out-performance easy to see. The value of the chart is less important than the slope of the line. If the slope is positive, the first asset is outperforming the second, and vice versa. Affectionately known as the Canto chart, it was used effectively in Canto (2006).

For distributional analysis, a few graphics may be useful. `chart.Boxplot` is an example of a graphic that is difficult to create in Excel and is under-utilized as a result. A boxplot of returns is, however, a very useful way to instantly observe the shape of hundreds or thousands of returns in a manner that makes them easy to compare to one another. `chart.Histogram` and `chart.QQPlot` are two charts from the FinMetrics packages with only small modifications.

Rolling performance is typically used as a way to assess stability of a return stream. Although perhaps it doesn't get much credence in the financial literature as it derives from work in digital signal processing, many practitioners find it a useful way to examine and segment performance and risk periods. See `chart.RollingPerformance`, which is a way to display different metrics over rolling time periods. `chart.RollingMean` is a specific example of a rolling mean and one standard deviation bands. A group of related metrics is offered in `charts.RollingPerformance`. These charts use utility functions such as `rollingFunction` and `rollingStat`. `chart.RollingCorrelation` shows how correlations change over rolling periods. `chart.RollingRegression` displays the coefficients of a linear model fitted over rolling periods. A group of charts in `charts.RollingRegression` displays alpha, beta, and R-squared estimates in three aligned charts in a single device.

There are also a number of charts or groups of charts that remain unfinished or were experimental and may be modified or removed in the future. `chart.RegressionDiagnostics`, `chart.Correlation`, and `chart.Correlation.color` are examples. The latter two are slight modifications of code from the package MASS and the R base. See their documentation for more information.

We have been greatly inspired by other peoples' work, some of which is on display at <http://addictedtor.free.fr/>. Particular inspiration came from Dirk Eddelbuettel and John Bollinger for their work at <http://addictedtor.free.fr/graphiques/RGraphGallery.php?graph=65>.

Risk Analysis

Many methods have been proposed to measure, monitor, and control the risks of a diversified portfolio. Perhaps a few definitions are in order on how different risks are generally classified. *Market Risk* is the risk to the portfolio from a decline in the market price of instruments in the portfolio. *Liquidity Risk* is the risk that the holder of an instrument will find that a position is illiquid, and will incur extra costs in unwinding the position resulting in a less favorable price for the instrument. In extreme cases of liquidity risk, the seller may be unable to find a buyer for the instrument at all, making the value unknowable or zero. *Credit Risk* is also sometimes called *Default Risk*, or the risk that promised payments on a loan or bond will not be made, or that a convertible instrument will not be converted in a timely manner or at all. There are also additional *Counterparty Risks* in manual or over the counter markets, and in many complex derivatives. Tools have evolved to measure all these different components of risk. Processes must be put into place inside a firm to monitor the changing risks in a portfolio, and to control the magnitude of risks. For an extensive treatment of these topics, see Litterman, Gumerlock, et. al.(1998).

The simplest risk measure in common use is volatility, usually modeled quantitatively with a univariate standard deviation on a portfolio. See `StdDev`. Volatility or Standard Deviation is an appropriate risk measure when the distribution of returns is normal or resembles a random walk, and may be annualized using `StdDev.annualized`. Many assets, including hedge funds, commodities, options, and even most common stocks over a sufficiently long period, do not follow a normal distribution. For such common but non-normally distributed assets, a more sophisticated approach than standard deviation/volatility is required to adequately model the risk.

Markowitz, in his Nobel acceptance speech and in several papers, proposed that `SemiVariance` would be a better measure of risk than variance. See Zin, Markowitz, Zhao(2006) http://papers.ssrn.com/sol3/papers.cfm?abstract_id=910640. This measure is also called `SemiDeviation`. The more general case of downside deviation is implemented in the function `DownsideDeviation`, as proposed by Sortino and Price(1994), where the minimum acceptable return (MAR) is a parameter to the function. It is interesting to note that variance and mean return can produce a smoothly elliptical efficient frontier for portfolio optimization utilizing `solve.QP` or `portfolio.optim` or `MarkowitzPortfolio`. Use of semivariance or many other risk measures will not necessarily create a smooth ellipse, causing significant additional difficulties for the portfolio manager trying to build an optimal portfolio. We'll leave a more complete treatment and implementation of portfolio optimization techniques for another time.

Another very widely used downside risk measures is analysis of drawdowns, or loss from peak value achieved. The simplest method is to check the `maxDrawdown`, as this will tell you the worst cumulative loss ever sustained by the asset. If you want to look at all the drawdowns, you can `findDrawdowns` and `sortDrawdowns` in order from worst/major to smallest/minor. The `UpDownRatios` function will give you some insight into the impacts of the skewness and kurtosis of the returns, and letting you know how length and magnitude of up or down moves compare to each other. You can also plot drawdowns with `chart.Drawdown`.

One of the most commonly used and cited measures of the risk/reward tradeoff of an investment or portfolio is the `SharpeRatio`, which measures return over standard deviation. If you are comparing multiple assets using Sharpe, you should use `SharpeRatio.annualized`. The `SortinoRatio` utilizes mean return over `DownsideDeviation` below the MAR as the risk measure to produce a similar ratio that is more sensitive to downside risk. Favre and Galeano(2002) propose utilizing the ratio of expected excess return over the Cornish-Fisher `modifiedVaR` to produce `SharpeRatio.modified`. `TreynorRatio` is also similar to the Sharpe Ratio, except it uses `CAPM.beta` in place of the volatility measure to produce the ratio of the investment's excess return over the beta.

One of the newest statistical methods developed for analyzing the risk of financial instruments is `Omega`. Omega analytically constructs a cumulative distribution function, in a manner similar

to `chart.QQPlot`, but then extracts additional information from the location and slope of the derived function at the point indicated by the risk quantile that the researcher is interested in. Omega seeks to combine a large amount of data about the shape, magnitude, and slope of the distribution into one method. Omega is still a very new method, and the academic literature is still exploring the best manner to utilize Omega in a risk measurement and control process, or in portfolio construction.

Any risk measure should be viewed with suspicion if there are not a large number of historical observations of returns for the asset in question available. Depending on the measure, how many observations are required will vary greatly from a statistical standpoint. As a heuristic rule, ideally you will have data available on how the instrument performed through several economic cycles and shocks. When such a long history is not available, the investor or researcher has several options. A full discussion of the various approaches is beyond the scope of this introduction, so we will merely touch on several areas that an interested party may wish to explore in additional detail. Examining the returns of assets with a similar style, industry, or asset class to which the asset in question is highly correlated and shares other characteristics can be quite informative. Factor analysis may be utilized to uncover specific risk factors where transparency is not available. Various resampling (see `tsbootstrap` and simulation methods are available in R to construct an artificially long distribution for testing. If you use a method such as Monte Carlo simulation or the bootstrap, it is often valuable to use `chart.Boxplot` to visualize the different estimates of the risk measure produced by the simulation, to see how small (or wide) a range the estimates cover, and thus gain a level of confidence with the results. Proceed with extreme caution when your historical data is lacking. Problems with lack of historical data are a major reason why many institutional investors will not invest in an alternative asset without several years of historical return data available.

Value at Risk - VaR

Traditional mean-VaR: In the early 90's, academic literature started referring to "value at risk", typically written as VaR. Take care to capitalize VaR in the commonly accepted manner, to avoid confusion with var (variance) and VAR (vector auto-regression). With a sufficiently large data set, you may choose to utilize a non-parametric VaR estimation method using the historical distribution and the probability quantile of the distribution calculated using `qnorm`. The negative return at the correct quantile (usually 95% or 99%), is the non-parametric VaR estimate. J.P. Morgan's RiskMetrics parametric mean-VaR was published in 1994 and this methodology for estimating parametric mean-VaR has become what people are generally referring to as "VaR" and what we have implemented as `VaR.traditional`. See *Return to RiskMetrics: Evolution of a Standard* <http://www.riskmetrics.com/r2rovv.html>, `fPortfolio` has also implemented traditional mean-VaR as the `VaR` function. Parametric traditional VaR does a better job of accounting for the tails of the distribution by more precisely estimating the tails below the risk quantile. It is still insufficient if the assets have a distribution that varies widely from normality.

The R package `VaR` contains methods for simulating and estimating lognormal `VaR.norm` and generalized Pareto `VaR.gpd` distributions to overcome some of the problems with nonparametric or parametric mean-VaR calculations on a limited sample size. There is also a `VaR.backtest` function to apply simulation methods to create a more robust estimate of the potential distribution of losses. The `VaR` package also provides plots for its functions.

Conditional VaR and Beyond VaR: The `fPortfolio` package has implemented Conditional Value at Risk, also called Expected Shortfall (not to be confused with shortfall probability, which is much less useful), in functions `CVaR` and `CVaRplus`. Expected Shortfall attempts to measure the magnitude of the average loss exceeding the traditional mean-VaR. Expected Shortfall has proven to be a reasonable risk predictor for many asset classes. See *Uryasev(2000)* and *Sherer and Martin(2005)* for more information on Conditional Value at Risk. A similar measure called Beyond VaR implemented here as `VaR.Beyond` also attempts to estimate average loss beyond traditional mean-VaR. Adding mean-VaR and Beyond VaR will produce a measure very similar to `CVaRplus`. Please note that your mileage will vary; expect that values obtained from the normal distribution may differ

radically from the real situation, depending on the assets under analysis. Both CVaR and Beyond-VaR have been shown to perform significantly better than mean-VaR with respect to non-normal distributions, especially when estimated using the bootstrap or Monte Carlo.

Modified Cornish-Fisher VaR: The limitations of traditional mean-VaR are all related to the use of a symmetrical distribution function. Use of simulations, resampling, or Pareto distributions all help in making a more accurate prediction, but they are still flawed for assets with significantly non-normal (skewed or kurtotic) distributions. *Huisman(1999)* and *Favre and Galleano(2002)* propose to overcome this extensively documented failing of traditional VaR by directly incorporating the higher moments of the return distribution into the VaR calculation. This new VaR measure incorporates skewness and kurtosis via an analytical estimation using a Cornish-Fisher (special case of a Taylor) expansion. The resulting measure is referred to variously as “Cornish-Fisher VaR” or “Modified VaR”. We provide this measure as function `VaR.CornishFisher` with equivalent alias `modifiedVaR`, Modified VaR produces the same results as traditional mean-VaR when the return distribution is normal, so it may be used as a direct replacement. Many papers in the finance literature have reached the conclusion that Modified VaR is a superior measure, and may be substituted in any case where mean-VaR would previously have been utilized.

Marginal, Incremental, and Component VaR: Marginal VaR is the difference between the VaR of the portfolio without the asset in question and the entire portfolio. The `VaR.Marginal` function calculates Marginal VaR for all instruments in the portfolio. Marginal VaR as provided here may use traditional mean-VaR or Modified VaR for the calculation. Per *Artzner, et.al.(1997)* properties of a coherent risk measure include subadditivity (risks of the portfolio should not exceed the sum of the risks of individual components) as a significantly desirable trait. VaR measures, including Marginal VaR, on individual components of a portfolio are *not* subadditive. *Sherer and Martin(2005)* p. 180 report that CVaR may be a subadditive measure under many circumstances, but present only a special rather than the general case. Clearly, a general subadditive risk measure for downside risk is required. Incremental or Component VaR attempt to provide just such a measure. In Incremental or Component VaR, the Component VaR value for each element of the portfolio will sum to the total VaR of the portfolio. Several EDHEC papers suggest utilizing Modified VaR instead of mean-VaR in the Incremental and Component VaR calculation. We have not yet succeeded in implementing an Incremental and Component VaR calculation that utilizes Modified Cornish-Fisher VaR. We hope to add it to a future version of this library.

Which VaR measure to use will depend greatly on the portfolio and instruments being analyzed. If there is any generalization to be made on VaR measures, this author will agree with *Bali and Gokcan(2004)* who conclude that “the VaR estimations based on the generalized Pareto distribution and the Cornish-Fisher approximation perform best”.

Wrapper and Utility Functions

R is a very powerful environment for manipulating data. It can also be quite confusing to a user more accustomed to Excel or even MatLAB. As such, we have written some wrapper functions that may aid you in coercing data into the correct forms or finding data that you need to use regularly. To simplify the management of multiple-source data stored in R in multiple data formats, we have provided `checkDataMatrix` and `checkDataVector`. These functions will attempt to coerce data in and out of R’s multitude of mostly fungible data classes into the class required for a particular analysis. In general, the use of these data-coercion functions has been hidden inside the business functions provided. They may also save you time and trouble in your own code and functions outside of the functionality provided by the PerformanceAnalytics library.

When you are analyzing relative or absolute performance of investments, you need to analyze returns, but much data is available only as prices. We have provided the simple wrapper function `CalculateReturns` to address this by taking a stream of prices and calculating simple or compounded returns from the price vector. The excellent `tseries` library includes the function `get.hist.quote` for retrieving market data from online sources. Many of the functions

in PerformanceAnalytics require either a benchmark or a risk free rate, and many practitioners will not have access to a formal database of returns. To facilitate the examples and provide an example of how to retrieve and coerce the data, we have provided functions for S&P 500 returns in the [download.SP500PriceReturns](#) and the 13-day US Treasury Bill in [download.RiskFree](#).

We also provide wrappers for cumulative maxima of the returns in a multicolumn series in [cummax.column](#) and for the cumulative product (compound return) in [cumprod.column](#).

Further Work

Optimization

Performance Attribution

Factor Analysis

Hedge Selection and Analysis

Shock/Scenario Analysis

We have attempted to standardize function parameters and variable names, but more work exists to be done here.

In some cases, there are separate functions for single column and multicolumn data, which should be combined into a single function for ease of use and clarity.

There are functions included in this package that would benefit from more complex examples.

There is also an opportunity to extend several of the co-moments and risk measures to use multivariate estimates of the higher moments, which should produce more robust estimates of the moments of a portfolio of multiple diverse assets.

Any comments, suggestions, or code patches to these ends are invited.

If you've implemented anything in the list above, please consider donating it for inclusion in a later version of this package.

Acknowledgements

Data series [edhec](#) used in PerformanceAnalytics and related publications with the kind permission of the EDHEC Risk and Asset Management Research Center.

http://www.edhec-risk.com/indexes/pure_style

Prototypes of the drawdowns functionality were provided by Sankalp Upadhyay, and modified with permission.

Thanks to Joe Byers and Dirk Eddelbuettel for comments on early versions of these functions.

Thanks to the R-SIG-Finance community without whom this library would not be possible.

Author(s)

Brian G. Peterson

Peter Carl

Maintainer: Brian G. Peterson <brian@braverock.com> <http://braverock.com/brian/R/>

References

Canto V. *Understanding Asset Allocation*. FT Prentice Hall. 2006.

Lhabitant F. *Hedge Funds: Quantitative Insights*. Wiley. 2004.

Litterman R., Gumerlock R., et. al. *The Practice of Risk Management: Implementing Processes for*

- Managing Firm-Wide Market Risk*. Euromoney. 1998.
- Murrel, Paul. *R Graphics*. Chapman and Hall. 2006.
- Ruppert, David. *Statistics and Finance, an Introduction*. Springer. 2004.
- Scherer, B. and Martin, D. *Modern Portfolio Optimization*. Springer. 2005.
- Tsay, R. *Analysis of Financial Time Series*. Wiley. 2001.
- Zivot, E. and Wang, Z. *Modeling Financial Time Series with S-Plus: second edition*. Springer. 2006.

See Also

CRAN task view on Empirical Finance <http://cran.r-project.org/src/contrib/Views/Econometrics.html>

Grant Farnsworth's Econometrics in R <http://cran.r-project.org/doc/contrib/Farnsworth-Econometrics.pdf>

Collection of R charts and graphs <http://addictedtor.free.fr/graphiques/>

Index

*Topic package

- PerformanceAnalytics-package,
1
- ActivePremium, 3
- aggregate, 1
- CalculateReturns, 1, 2, 8
- CAPM.alpha, 3
- CAPM.beta, 3, 6
- CAPM.CML, 3
- CAPM.CML.slope, 3
- CAPM.RiskPremium, 3
- CAPM.SML.slope, 3
- chart.Bar, 5
- chart.BarVaR, 5
- chart.Boxplot, 5, 7
- chart.Correlation, 5
- chart.Correlation.color, 5
- chart.CumReturns, 5
- chart.Drawdown, 5, 6
- chart.Histogram, 3, 5
- chart.QQPlot, 3, 5, 7
- chart.RegressionDiagnostics, 5
- chart.RelativePerformance, 5
- chart.RiskReturnScatter, 5
- chart.RollingCorrelation, 5
- chart.RollingMean, 5
- chart.RollingPerformance, 5
- chart.RollingRegression, 5
- chart.Scatter, 5
- charts.PerformanceSummary, 5
- charts.RollingPerformance, 5
- charts.RollingRegression, 5
- checkDataMatrix, 8
- checkDataVector, 8
- CoKurtosis, 4
- CoSkewness, 4
- cummax.column, 9
- cumprod.column, 9
- CVaR, 7
- CVaRplus, 7
- download.RiskFree, 9
- download.SP500PriceReturns, 9
- DownsideDeviation, 6
- edhec, 9
- factanal, 4
- findDrawdowns, 6
- get.hist.quote, 8
- InformationRatio, 3
- KellyRatio, 3
- kurtosis, 4
- MarkowitzPortfolio, 6
- max, 2
- maxDrawdown, 6
- mean, 2
- mean.geometric, 2
- mean.LCL, 2
- mean.stderr, 2
- mean.UCL, 2
- min, 2
- modifiedVaR, 6, 8
- Omega, 6
- PerformanceAnalytics
(PerformanceAnalytics-package),
1
- PerformanceAnalytics-package, 1
- pfolioReturn, 2
- portfolio.optim, 6
- princomp, 4
- qnorm, 7
- quantile, 2
- range, 2
- rcorr, 3
- Return.annualized, 2
- rollingFunction, 5
- rollingStat, 5
- SemiDeviation, 6

SemiVariance, 3, 6
SharpeRatio, 3, 6
SharpeRatio.annualized, 2, 6
SharpeRatio.modified, 6
skewness, 3, 4
solve.QP, 6
sortDrawdowns, 6
SortinoRatio, 3, 6
StdDev, 3, 6
StdDev.annualized, 2, 6
SystematicKurtosis, 4
SystematicSkewness, 4

table.AnnualizedReturns, 5
table.CAPM, 5
table.Correlation, 5
table.Downsiderisk, 5
table.HigherMoments, 5
table.MonthlyReturns, 2, 5
table>Returns, 4
table.RollingPeriods, 5
TrackingError, 3
TreynerRatio, 6
tsbootstrap, 7

UpDownRatios, 6

VaR, 7
var, 3
VaR.backtest, 7
VaR.Beyond, 7
VaR.CornishFisher, 8
VaR.gpd, 7
VaR.Marginal, 8
VaR.norm, 7
VaR.traditional, 7