

PerformanceAnalytics Charts and Tables Overview

Peter Carl & Brian G. Peterson

December 29, 2007

Abstract

This vignette gives a brief overview of (some of) the graphics and display wrapper functionality contained in **PerformanceAnalytics** including most of the charts and tables . For a more complete overview of the package's functionality and extensibility see the manual pages. We develop the examples using data for six (hypothetical) managers, a peer index, and an asset class index.

Contents

1	Introduction	2
2	Set up PerformanceAnalytics	2
2.1	Install PerformanceAnalytics	2
2.2	Load and review data	3
2.3	Create charts and tables for presentation	5
3	Create Charts and Tables	8
3.1	Create performance charts	8
3.2	Create a monthly returns table	10
3.3	Calculate monthly statistics	10
3.4	Compare distributions	11
3.5	Show relative return and risk	12
3.6	Examine performance consistency	16
3.7	Display relative performance	16
3.8	Measure relative performance to a benchmark	16
3.9	Calculate Downside Risk	21
4	Conclusion	24

1 Introduction

PerformanceAnalytics is a library of functions designed for evaluating the performance and risk characteristics of financial assets or funds. In particular, we have focused on functions that have appeared in the academic literature over the past several years, but had no functional equivalent in R.

Our goal for **PerformanceAnalytics** is to make it simple for someone to ask and answer questions about performance and risk as part of a broader investment decision-making process. There is no magic bullet here – there won’t be one *right* answer delivered in these metrics and charts. Investments must be made in context of investment objectives. But what this library aspires to do is *help* the decision-maker accrete evidence organized to answer a specific question that is pertinent to the decision at hand. Our hope is that using such tools to uncover information and ask better questions will, in turn, create a more informed investor and help them ask better quality decisions.

This vignette provides a demonstration of some of the capabilities of **PerformanceAnalytics**. We focus on the graphs and tables, but comment on some other metrics along the way. These examples are not intended to be complete, but they should provide an indication of the kinds of analysis that can be done. Other examples are available in the help pages of the functions described in the main page of **PerformanceAnalytics**.

2 Set up PerformanceAnalytics

These examples assume the reader has basic knowledge of R and understands how to install R, read and manipulate data, and create basic calculations. For further assistance, please see <http://cran.r-project.org/doc/manuals/R-intro.pdf> and other available materials at <http://cran.r-project.org>. This section will begin with installation, discuss the example data set, and provide an overview of charts attributes that will be used in the examples that follow.

2.1 Install PerformanceAnalytics

As of version 0.9.4, **PerformanceAnalytics** is available via CRAN, but you will need version 0.9.5 or later to replicate the examples that follow. R users with connectivity can simply type: `install.packages("PerformanceAnalytics")`. Or, if you already have an earlier version installed, see `update.packages`. A number of packages are required, including **Hmisc**, **zoo**, and the various **Rmetrics** packages such as **fBasics**, **fCalendar**, and **fExtremes**. After installing **PerformanceAnalytics**, load it into your active R session using `library("PerformanceAnalytics")`.

Figure 1: First Lines of the managers Object

```
> data(managers)
> head(managers)
```

	HAM1	HAM2	HAM3	HAM4	HAM5	HAM6	EDHEC.LS.EQ	SP500.TR	US.10Y.TR
Jan 1996	0.0100	NA	0.0359	0.0208	NA	NA	NA	0.0340	0.00380
Feb 1996	0.0215	NA	0.0295	0.0231	NA	NA	NA	0.0093	-0.03532
Mar 1996	0.0226	NA	0.0253	-0.0053	NA	NA	NA	0.0096	-0.01057
Apr 1996	0.0008	NA	0.0478	0.0200	NA	NA	NA	0.0147	-0.01739
May 1996	0.0158	NA	0.0337	0.0122	NA	NA	NA	0.0258	-0.00543
Jun 1996	-0.0086	NA	-0.0293	-0.0089	NA	NA	NA	0.0038	0.01507

	US.3m.TR
Jan 1996	0.00456
Feb 1996	0.00398
Mar 1996	0.00371
Apr 1996	0.00428
May 1996	0.00443
Jun 1996	0.00412

2.2 Load and review data

First we load the data used in all of the examples that follow. As you can see in Figure 1, `managers` is a data frame that contains columns of monthly returns for six hypothetical asset managers (HAM1 through HAM6), the EDHEC Long-Short Equity hedge fund index, the S&P 500 total returns, and total return series for the US Treasury 10-year bond and 3-month bill. Monthly returns for all series end in December 2006 and begin at different periods starting from January 1996. Similar data could be constructed using `mymanagers=read.csv("/path/to/file/mymanagers.csv", row.names=1)`, where the first column contains dates in the YYYY-MM-DD format.

A quick sidenote: this library is applicable to return (rather than price) data, and has been tested mostly on a monthly scale. Many library functions will work with regular data at different scales (e.g., daily, weekly, etc.) or irregular return data as well. See function `CalculateReturns` for calculating returns from prices, and be aware that the `zoo` library's `aggregate` function has methods for `tseries` and `zoo` timeseries data classes to rationally coerce irregular data into regular data of the correct periodicity.

With the data object in hand, we group together columns of interest make the examples easier to follow. As shown in Figure 2, we first confirm the dimensions of the object and the overall length of the timeseries.

Figure 2: Assess and Organize the Data

```
> dim(managers)

[1] 132  10

> managers.length = dim(managers)[1]
> colnames(managers)

[1] "HAM1"      "HAM2"      "HAM3"      "HAM4"      "HAM5"
[6] "HAM6"      "EDHEC.LS.EQ" "SP500.TR"  "US.10Y.TR" "US.3m.TR"

> manager.col = 1
> peers.cols = c(2, 3, 4, 5, 6)
> indexes.cols = c(7, 8)
> Rf.col = 10
> trailing12.rows = ((managers.length - 11):managers.length)
> trailing12.rows

[1] 121 122 123 124 125 126 127 128 129 130 131 132

> trailing36.rows = ((managers.length - 35):managers.length)
> trailing60.rows = ((managers.length - 59):managers.length)
> frInception.rows = (length(managers[, 1]) - length(managers[,
+ 1][!is.na(managers[, 1])]) + 1):length(managers[, 1])
```

Next, we group the columns together. The first column, HAM1, contains the subject of our analysis in the examples below. Columns two through six contain returns of other managers using a similar investing style that might be considered substitutes – a “peer group,” of sorts. Columns seven and eight contain key indexes. The EDHEC Long/Short Equity index is a “peer index,” the S&P 500 Total Return index is an “asset-class index”. We combine those into a set of “indexes” for later analysis. Column nine we skip for the time being; column 10 we use as the risk free rate for each month.

Then we do the same thing for the rows of interest. We calculate the row numbers that represent different trailing periods and keep those handy for doing comparative analyses.

2.3 Create charts and tables for presentation

Graphs and charts help to organize information visually. Our goals in creating these functions were to simplify the process of creating well-formatted charts that are frequently used for portfolio analysis and to create print-quality graphics that may be used in documents for broader consumption. R’s graphics capabilities are substantial, but the simplicity of the output of R’s default graphics functions such as `plot` does not always compare well against graphics delivered with commercial asset or portfolio analysis software from places such as MorningStar or PerTrac.

Color Palettes

We have set up some specific color palattes designed to create readable line and bar graphs with specific objectives. We use this approach (rather than generating them on the fly) for two reasons: first, there are fewer dependencies on libraries that don’t need to be called dynamically; and second, to guarantee the color used for the n-th column of data. Figure 3 shows some examples of the different palates.

The first category of colorsets are designed to provide focus to the data graphed as the first element, and include `redfocus`, `bluefocus`, and `greenfocus`. These palettes are best used when there is an important data set for the viewer to focus on. The other data provide some context, so they are graphed in diminishing values of gray. These were generated with `RColorBrewer`, using the 8 level “grays” palette and replacing the darkest gray with the focus color. To coordinate these colorsets with the equal-weighted colors below, replace the highlight color with the first color of the equal weighted palette from below. This will coordinate charts with different purposes.

The second category is sets of equal-weighted colors. These colorsets are useful for when all of the data should be observed and distinguishable on a line graph. The different numbers in the name indicate the number of colors generated (six colors is probably the maximum for a readable linegraph, but we provide as many as twelve). Examples include `rainbow12equal` through `rainbow4equal`, in steps of two. These colorsets were

generated with `rainbow(12, s = 0.6, v = 0.75)`. The `rich12equal` and other corresponding colorsets were generated with with package `gplots` function `rich.colors(12)`. Similarly `tim12equal` and similar colorsets were generated with with package `fields` function `tim.colors(12)`, a function said to emulate the Matlab colorset. The `dark8equal`, `dark6equal`, `set8equal`, and `set6equal` colorsets were created with package `RColorBrewer`, e.g., `brewer.pal(8, "Dark2")`. A third category is a set of monochrome colorsets, including `greenmono`, `bluemono`, `redmono`, and `gray8mono` and `gray6mono`.

To see what these lists contain, just type the name.

```
> tim12equal
```

```
[1] "#00008F" "#0000EA" "#0047FF" "#00A2FF" "#00FEFF" "#5AFFA5" "#B5FF4A"  
[8] "#FFED00" "#FF9200" "#FF3700" "#DB0000" "#800000"
```

These are just lists of strings that contain the RGB codes of each color. You can easily create your own if you have a particular palate that you like to use. Alternatively, you can use the R default colors. For more information, see:

[HTTP://research.stowers-institute.org/efg/R/Color/Chart/index.htm](http://research.stowers-institute.org/efg/R/Color/Chart/index.htm)

Symbols

Similarly, there are a few sets of grouped symbols for scatter charts . These include `opensymbols`, `closedsymbols`, `fillsymbols`, `linesymbols`, and `allsymbols`.

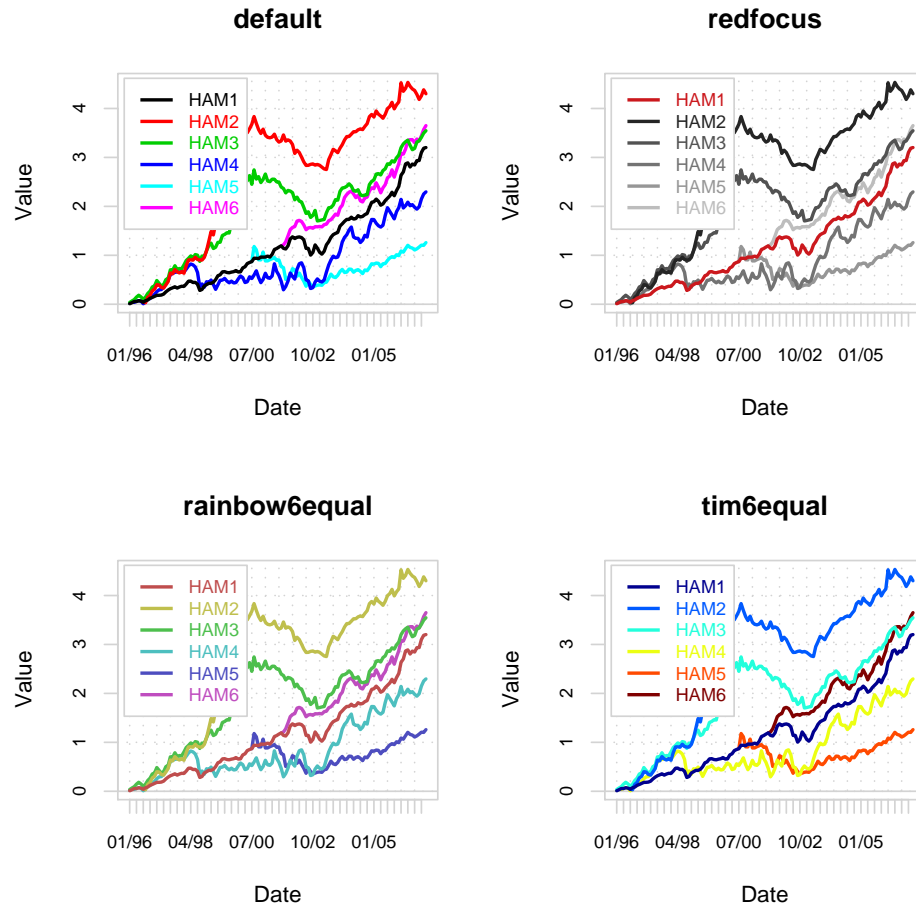
Legend locations

In the single charts the legend can be moved around on the plot. There are nine locations that can be specified by keyword: "bottomright", "bottom", "bottomleft", "left", "topleft", "top", "topright", "right" and "center". This places the legend on the inside of the plot frame at the given location. Further information can be found in `xy.coord`. Most compound charts have fixed legend locations.

Other Parameters

We have tried to leave access to all of the extensive list of parameters available in R's traditional graphics. For more information, see `plot.default` and `par`. In the example above, we passed `lwd` the value of 2, which affects the line width. We might also alter the line type or other parameter.

Figure 3: Examples of Color Palates



3 Create Charts and Tables

With that, we are ready to analyze the sample data set. This section starts with a set of charts that provide a performance overview, some tables for presenting the data and basic statistics, and then discusses ways to compare distributions, relative performance, and downside risk.

3.1 Create performance charts

Figure 4 shows a three-panel performance summary chart and the code used to generate it. The top chart is a normal cumulative return or wealth index chart that shows the cumulative returns through time for each column. For data with later starting times, set the parameter `begin = "axis"`, which starts the wealth index of all columns at 1, or `begin = "first"`, which starts the wealth index of each column at the wealth index value attained by the first column of data specified. The second of these settings, which is the default, allows the reader to see how the two indexes would compare had they started at the same time regardless of the starting period. In addition, the y-axis can be set to a logarithmic value so that growth can be compared over long periods. This chart can be generated independently using `chart.CumReturns`.

The second chart shows the individual monthly returns overlaid with a rolling measure of tail risk referred to as Cornish Fisher Value-at-Risk (VaR) or Modified VaR. Alternative risk measures, including standard deviation (specified as `"StdDev"`) and traditional Value-at-Risk (`"VaR"`), can be specified using the `method` parameter. Note that `StdDev` and `VaR` are symmetric calculations, so a high and low measure will be plotted. `ModifiedVaR`, on the other hand, is asymmetric and only a lower bound is drawn. These risk calculations are made on a rolling basis from inception, or can be calculated on a rolling window by setting `width` to a value of the number of periods. These calculations should help the reader to identify events or periods when estimates of tail risk may have changed suddenly, or to help evaluate whether the assumptions underlying the calculation seem to hold. The risk calculations can be generated for all of the columns provided by using the `all` parameter. When set to `TRUE`, the function calculates risk lines for each column given and may help the reader assess relative risk levels through time. This chart can be generated using `chart.BarVaR`.

The third chart in the series is a drawdown or underwater chart, which shows the level of losses from the last value of peak equity attained. Any time the cumulative returns dips below the maximum cumulative returns, it's a drawdown. This chart helps the reader assess the synchronicity of the loss periods and their comparative severity. As you might expect, this chart can also be created using `chart.Drawdown`.

Figure 4: Draw a Performance Summary Chart

```
> charts.PerformanceSummary(managers[, c(manager.col, indexes.cols)],
+   colorset = rich6equal, lwd = 2, ylog = TRUE)
```

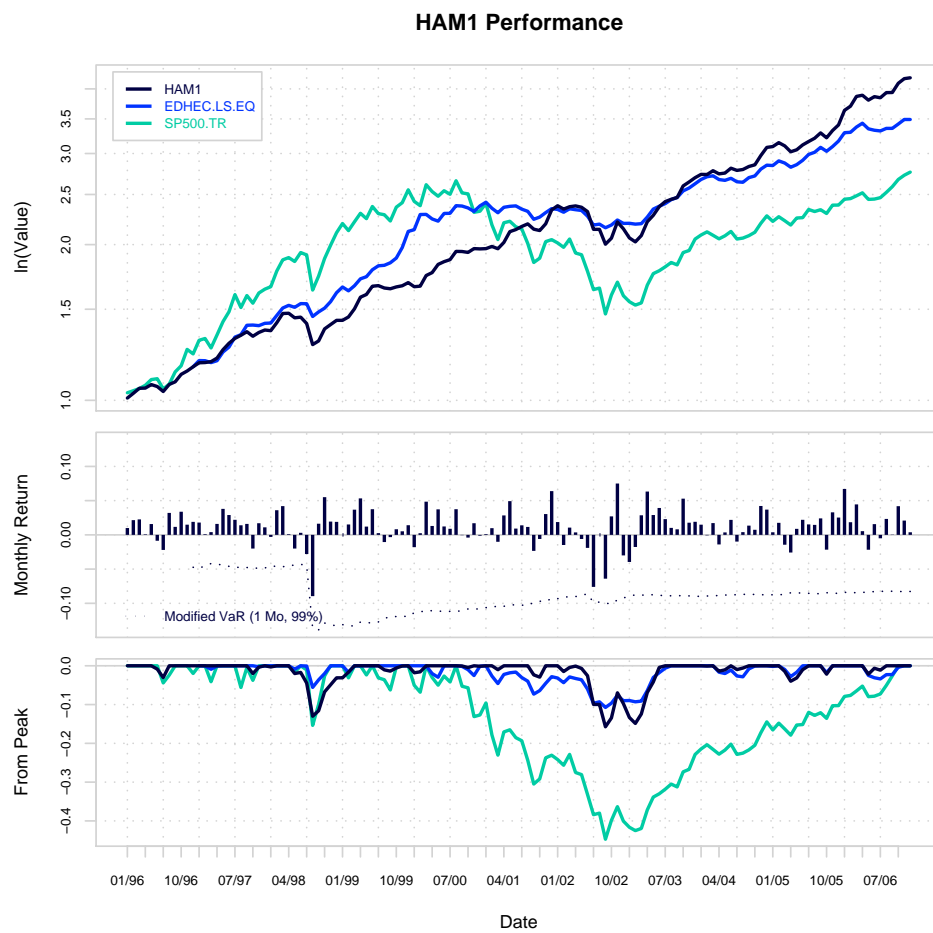


Figure 5: Create a Table of Calendar Returns

```
> t(table.CalendarReturns(managers[, c(manager.col, indexes.cols)]))
```

	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006
Jan	1.0	1.8	-0.3	0.0	-1.8	0.1	1.9	-4.0	1.5	0.4	6.7
Feb	2.1	0.1	3.6	1.5	0.2	1.0	-1.5	-1.8	-0.1	1.8	1.8
Mar	2.3	0.4	4.2	3.7	4.9	-1.0	1.1	2.9	1.7	-1.4	4.5
Apr	0.1	1.6	0.1	5.3	1.3	2.8	0.4	6.3	-1.4	-2.6	0.5
May	1.6	3.8	-2.0	1.2	3.7	4.9	-0.6	2.9	0.4	0.9	-2.2
Jun	-0.9	2.9	0.3	3.8	1.2	0.9	-1.9	3.9	2.2	2.2	1.6
Jul	-2.2	2.2	-2.8	0.2	0.9	1.4	-7.6	2.3	-1.0	1.5	-0.5
Aug	3.2	1.4	-8.9	-1.1	3.8	1.2	0.0	1.0	0.4	1.5	2.3
Sep	1.2	1.6	1.6	-0.3	0.0	-2.3	-6.4	0.8	1.4	2.4	0.0
Oct	3.4	-2.0	5.5	0.8	-0.4	-0.6	2.7	5.3	0.7	-2.2	4.2
Nov	1.5	1.7	1.9	0.5	1.7	3.0	7.5	1.8	4.2	3.3	2.1
Dec	1.9	1.1	1.9	1.4	-0.1	6.4	-3.0	1.9	3.7	2.5	0.4
HAM1	16.1	17.8	4.4	18.3	16.2	18.9	-8.1	25.5	14.4	10.5	23.3
EDHEC.LS.EQ	NA	21.4	14.6	31.4	12.0	-1.2	-6.4	19.3	8.6	11.3	10.1
SP500.TR	23.0	33.4	28.6	21.0	-9.1	-11.9	-22.1	28.7	10.9	4.9	15.8

3.2 Create a monthly returns table

Summary statistics are then the necessary aggregation and reduction of (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 accumulated by calendar year is usually presented as a table, such as in `table.CalendarReturns`.

Figure 5 shows a table of returns formatted with years in rows, months in columns, and a total column in the last column. For additional columns, the annual returns will be appended as columns. Adding benchmarks or peers alongside the annualized data is helpful for comparing returns across calendar years. Because there are a number of columns in the example, we make the output easier to read by using the `t` function to transpose the resulting data frame.

3.3 Calculate monthly statistics

Likewise, the monthly returns statistics table in Figure 6 was created as a way to display a set of related measures together for comparison across a set of instruments or funds. In

Figure 6: Create a Table of Statistics

```
> table.MonthlyReturns(managers[, c(manager.col, peers.cols)])
```

	HAM1	HAM2	HAM3	HAM4	HAM5	HAM6
Observations	132.0000	125.0000	132.0000	132.0000	77.0000	64.0000
NAs	0.0000	7.0000	0.0000	0.0000	55.0000	68.0000
Minimum	-0.0895	-0.0429	-0.0738	-0.1800	-0.1386	-0.0402
Quartile 1	0.0000	-0.0105	-0.0066	-0.0213	-0.0184	-0.0034
Median	0.0132	0.0060	0.0107	0.0139	0.0045	0.0146
Arithmetic Mean	0.0112	0.0138	0.0122	0.0105	0.0034	0.0121
Geometric Mean	0.0109	0.0131	0.0115	0.0091	0.0025	0.0118
Quartile 3	0.0231	0.0248	0.0312	0.0440	0.0298	0.0276
Maximum	0.0750	0.1521	0.1774	0.1583	0.1660	0.0544
SE Mean	0.0022	0.0033	0.0032	0.0047	0.0051	0.0030
LCL Mean (0.95)	0.0069	0.0072	0.0058	0.0013	-0.0067	0.0062
UCL Mean (0.95)	0.0156	0.0203	0.0186	0.0197	0.0136	0.0180
Variance	0.0006	0.0014	0.0014	0.0029	0.0020	0.0006
Stdev	0.0251	0.0369	0.0371	0.0536	0.0447	0.0238
Skewness	-0.6871	1.4564	0.8091	-0.4198	-0.0131	-0.2312
Kurtosis	2.4001	2.4099	2.3632	0.8703	2.1288	-0.5305

this example, we are looking at performance “from inception” for most of the managers, so careful consideration needs to be given to missing data or unequal time series when interpreting the results. Most people will prefer to see such statistics for common or similar periods. Each of the individual functions can be called individually, as well.

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.

3.4 Compare distributions

For distributional analysis, a few graphics may be useful. The result of `chart.Boxplot`, shown in Figure 7, 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 observe the shape of large collections of asset returns in a manner that makes them easy to compare

to one another.

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. Figure 8 shows a histogram generated for HAM1 with different display options.

Look back at the results generated by `table.MonthlyReturns`. Differences between `var` and `SemiVariance` will help you identify `[fBasics]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 `[fBasics]kurtosis` function will by default return what is referred to as “excess kurtosis”, where zero 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.

3.5 Show relative return and risk

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, such as in Figure 9. Examples are in `Return.annualized`, `StdDev.annualized`, and `SharpeRatio.annualized`.

`chart.Scatter` is a utility scatter chart with some additional attributes that are used in `chart.RiskReturnScatter`. Different risk parameters may be used. The parameter `method` may be any of "modVaR", "VaR", or "StdDev".

Additional information can be overlaid, as well. If `add.sharpe` is set to a value, say `c(1,2,3)`, the function overlays Sharpe ratio line that indicates Sharpe ratio levels of one through three. Lines are drawn with a y-intercept of the risk free rate (`rf`) and the slope of the appropriate Sharpe ratio level. Lines should be removed where not appropriate (e.g., `sharpe.ratio = NULL`). With a large number of assets (or columns), the names may get in the way. To remove them, set `add.names = NULL`. A box plot may be added to the margins to help identify the relative performance quartile by setting `add.boxplots = TRUE`.

Figure 7: Create a Boxplot

```
> chart.Boxplot(managers[trailing36.rows, c(manager.col, peers.cols,  
+      indexes.cols)], main = "Trailing 36-Month Returns")
```

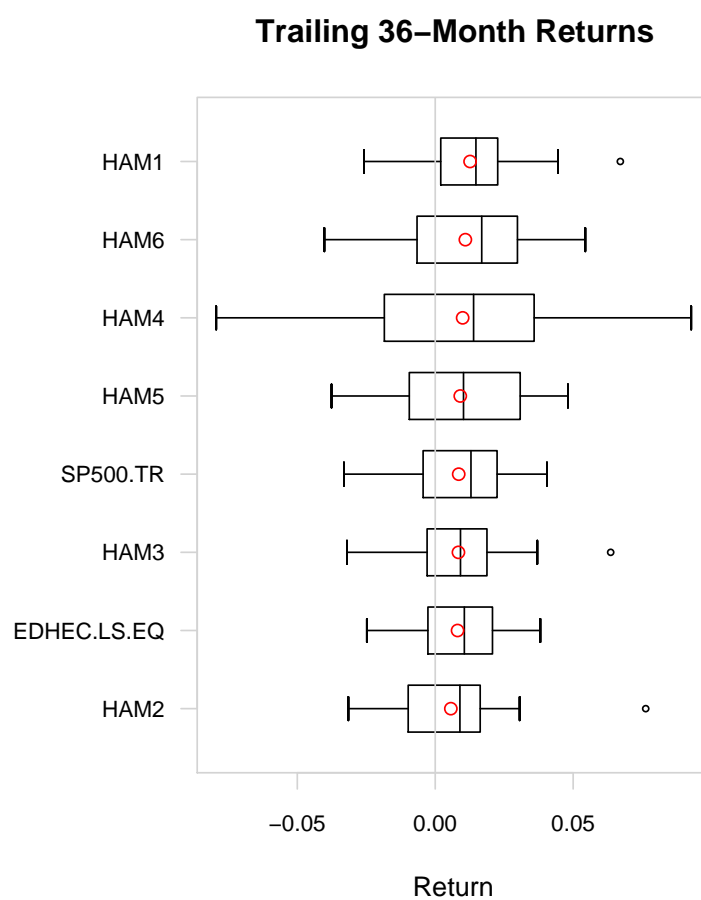


Figure 8: Create a Histogram of Returns

```
> layout(rbind(c(1, 2), c(3, 4)))
> chart.Histogram(managers[, 1, drop = F], main = "Plain", methods = NULL)
> chart.Histogram(managers[, 1, drop = F], main = "Density", breaks = 40,
+   methods = c("add.density", "add.normal"))
> chart.Histogram(managers[, 1, drop = F], main = "Skew and Kurt",
+   methods = c("add.centered", "add.rug"))
> chart.Histogram(managers[, 1, drop = F], main = "Risk Measures",
+   methods = c("add.risk"))
```

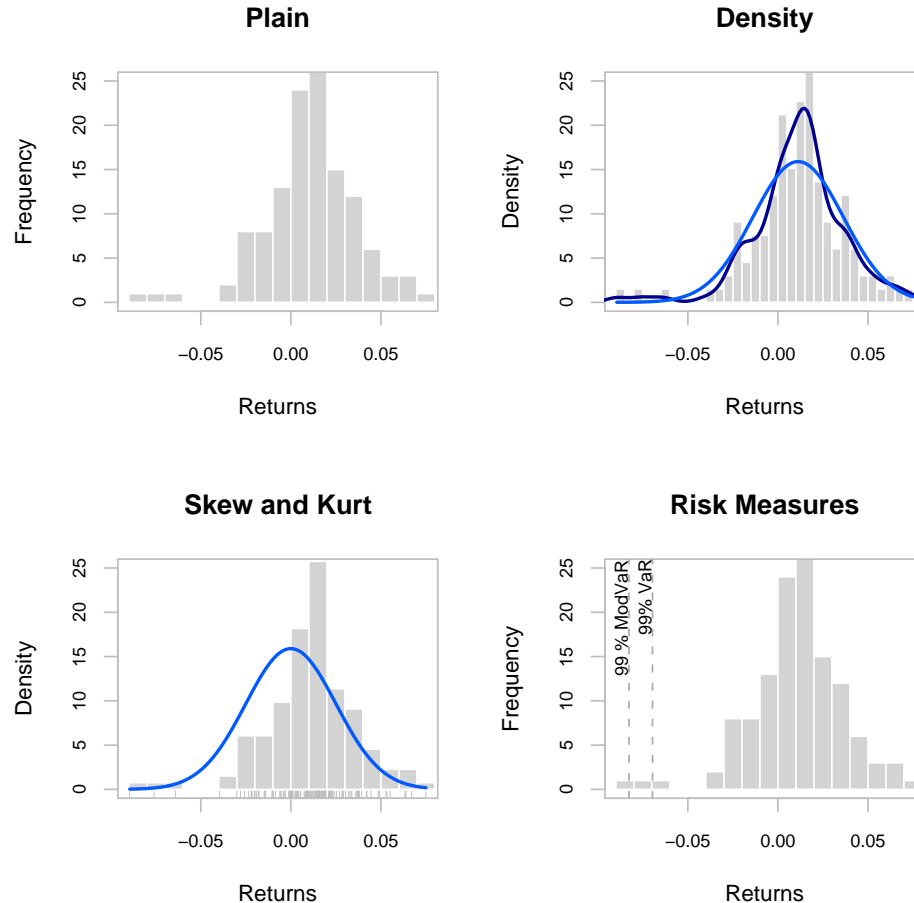
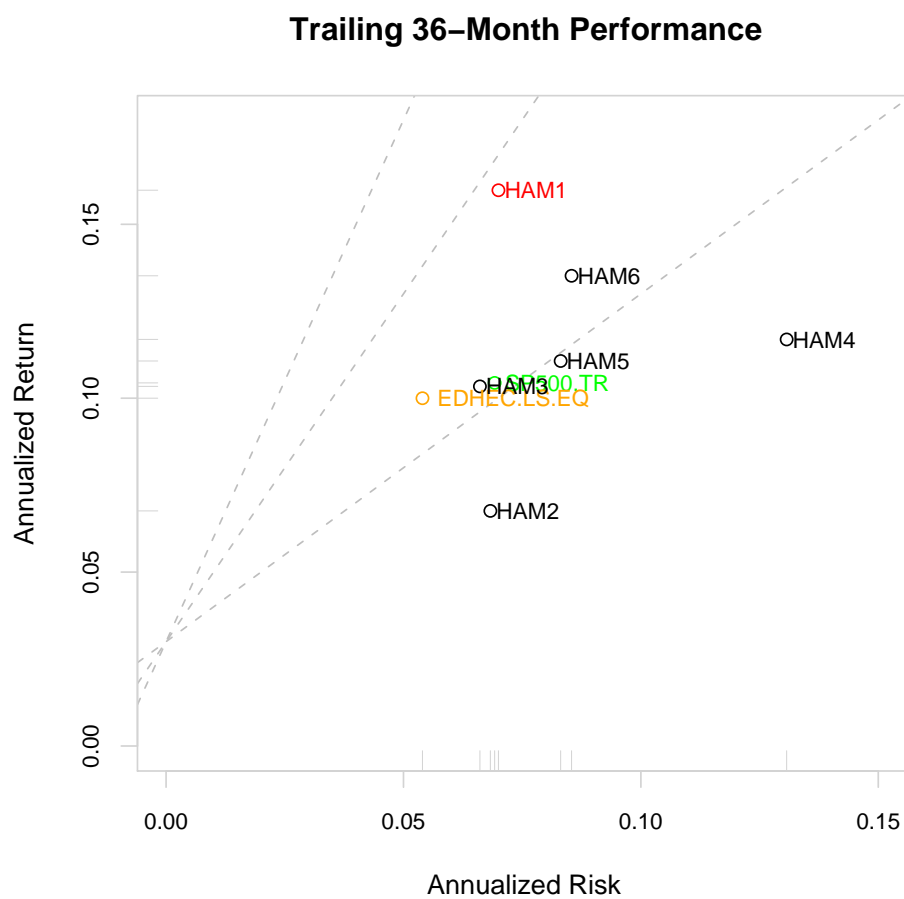


Figure 9: Show Relative Risk and Return

```
> chart.RiskReturnScatter(managers[trailing36.rows, 1:8], rf = 0.03/12,
+   main = "Trailing 36-Month Performance", colorset = c("red",
+   rep("black", 5), "orange", "green"))
```



3.6 Examine performance consistency

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 because of its roots in digital signal processing, many practitioners find rolling performance to be 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.

Figure 10 shows three panels, the first for rolling returns, the second for rolling standard deviation, and the third for rolling Sharpe ratio. These three panels each call `chart.RollingPerformance` with a different FUN argument, allowing any function to be viewed over a rolling window.

3.7 Display relative performance

The function `chart.RelativePerformance` shows the ratio of the cumulative performance for two assets 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 (numerator) is outperforming the second, and vice versa. Figure 11 shows the returns of the manager in question relative to each member of the peer group and the peer group index.

Looking at the data another way, we use the same chart to assess the peers individually against the asset class index. Figure 12 shows the peer group members' returns relative to the S&P 500. Several questions might arise: Who beats the S&P? Can they do it consistently? Are there cycles when they under-perform or outperform as a group? Which manager has outperformed the most?

3.8 Measure relative performance to a benchmark

Identifying and using a benchmark can help us assess and explain how well we are meeting our investment objectives, in terms of a widely held substitute. A benchmark can help us explain how the portfolios are managed, assess the risk taken and the return desired, and check that the objectives were respected. Benchmarks are used to get better control of the investment management process and to suggest ways to improve selection.

Modern Portfolio Theory is a collection of tools and techniques by which a risk-averse investor may construct an optimal portfolio. It encompasses the Capital Asset Pricing Model (CAPM), the efficient market hypothesis, and all forms of quantitative portfolio construction and optimization. CAPM 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 performance premium provided by an investment over a passive strategy (the benchmark) is provided by `ActivePremium`, which is the investment's annualized return

Figure 10: Examine Rolling Performance

```
> charts.RollingPerformance(managers[, c(manager.col, peers.cols,  
+   indexes.cols)], rf = 0.03/12, colorset = c("red", rep("darkgray",  
+   5), "orange", "green"), lwd = 2)
```

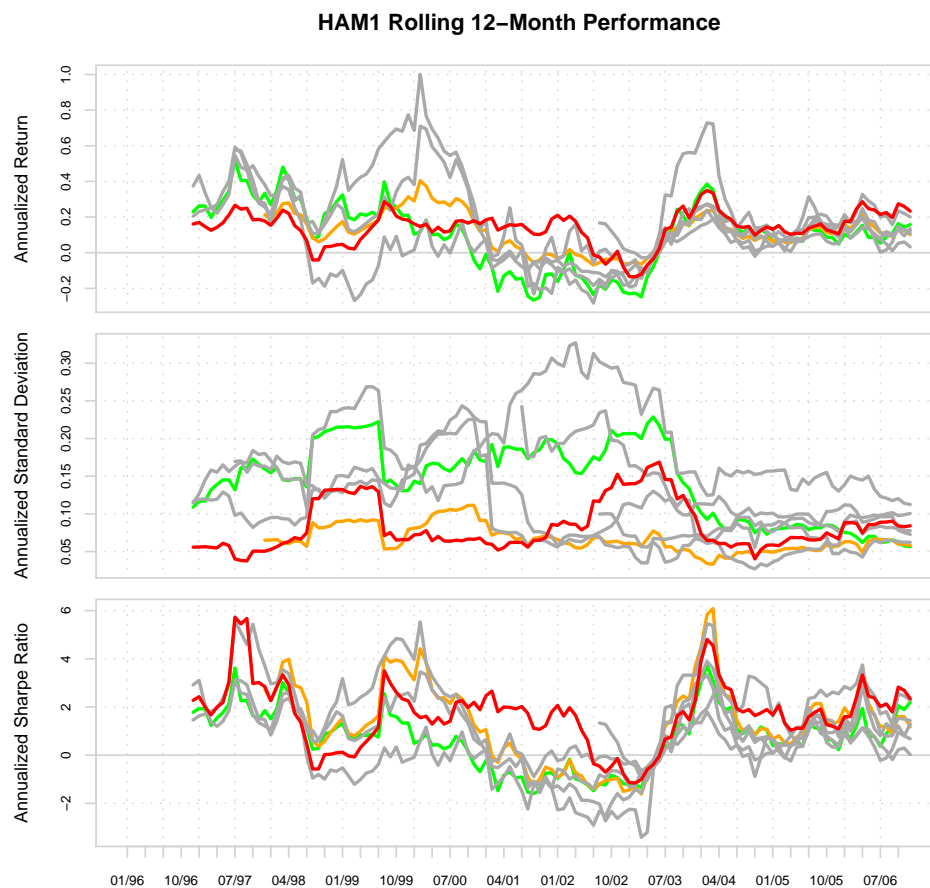


Figure 11: Examine Relative Performance of Assets

```
> chart.RelativePerformance(managers[, manager.col, drop = FALSE],  
+   managers[, c(peers.cols, 7)], colorset = tim8equal[-1], lwd = 2,  
+   legend.loc = "topleft")
```

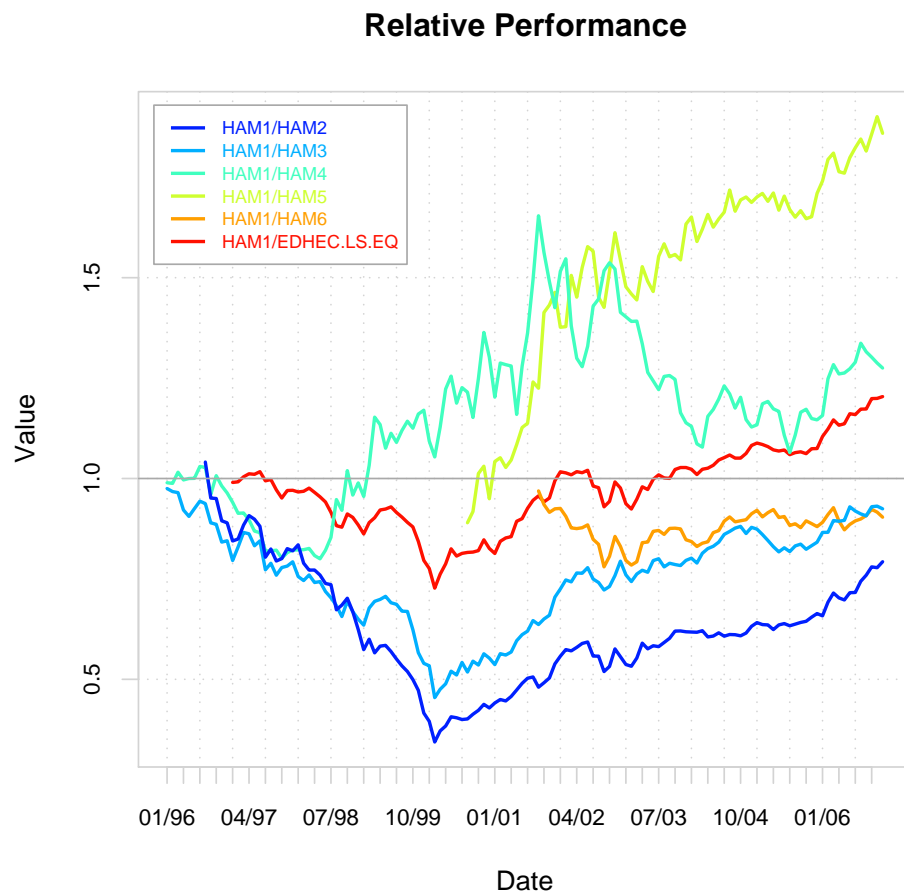


Figure 12: Examine Performance Relative to a Benchmark

```
> chart.RelativePerformance(managers[, c(manager.col, peers.cols)],  
+   managers[, 8, drop = F], colorset = rainbow8equal, lwd = 2,  
+   legend.loc = "topleft")
```

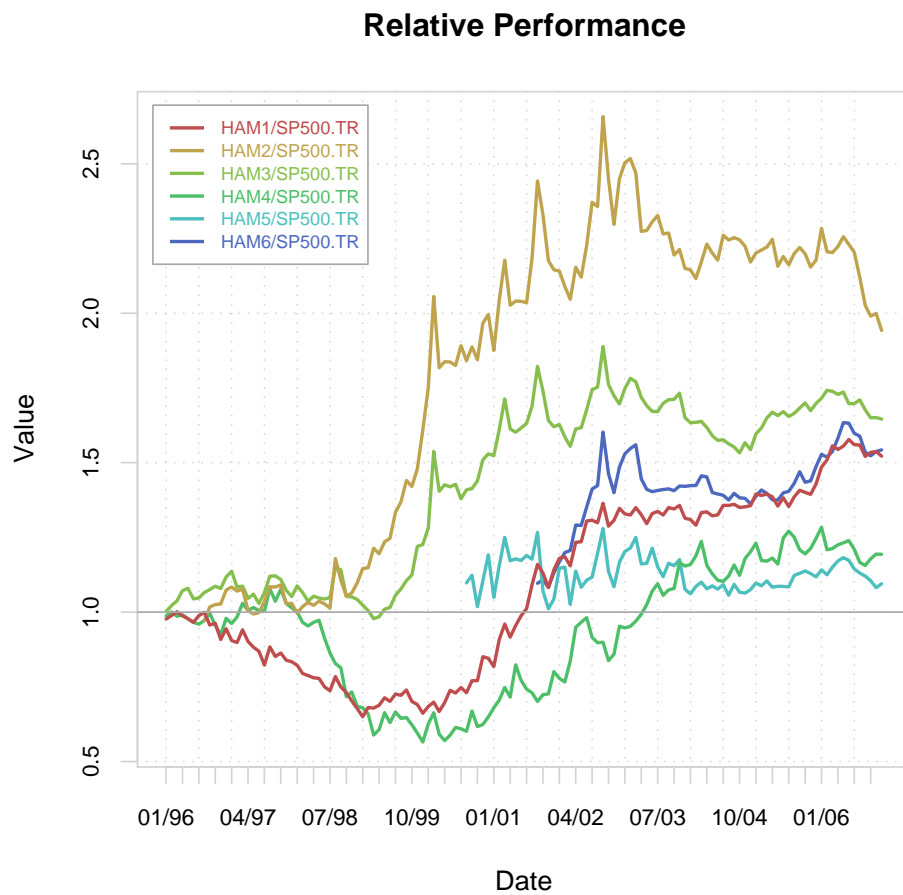


Figure 13: Create a Table of CAPM-Related Measures

```
> table.CAPM(managers[trailing36.rows, c(manager.col, peers.cols)],
+ managers[trailing36.rows, 8, drop = FALSE], rf = managers[trailing36.rows,
+ Rf.col, drop = FALSE])
```

	HAM1 to SP500.TR	HAM2 to SP500.TR	HAM3 to SP500.TR
Alpha	0.0061	0.0006	0.0015
Beta	0.6713	0.4178	0.7349
R-squared	0.4397	0.1715	0.5907
Annualized Alpha	0.0755	0.0076	0.0180
Correlation	0.6631	0.4142	0.7686
Correlation p-value	0.0000	0.0120	0.0000
Tracking Error	0.0868	0.0601	0.0021
Active Premium	0.0538	-0.0359	-0.0010
Information Ratio	0.6201	-0.5974	-0.4973
Treynor Ratio	0.1870	0.0857	0.0962
	HAM4 to SP500.TR	HAM5 to SP500.TR	HAM6 to SP500.TR
Alpha	0.0005	0.0015	0.0033
Beta	1.1570	0.8442	0.8574
R-squared	0.3697	0.4887	0.4830
Annualized Alpha	0.0059	0.0181	0.0399
Correlation	0.6080	0.6991	0.6950
Correlation p-value	0.0001	0.0000	0.0000
Tracking Error	0.0302	0.0119	0.0508
Active Premium	0.0120	0.0061	0.0299
Information Ratio	0.3984	0.5148	0.5889
Treynor Ratio	0.0724	0.0922	0.1186

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 **ActivePremium** divided by the **TrackingError**. **InformationRatio** may be used to rank investments in a relative fashion.

The code in Figure 13 creates a table of CAPM-related statistics that we can review and compare across managers. Note that we focus on the trailing-36 month period. There are, in addition to those listed, a wide variety of other CAPM-related metrics available. 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.

In a similar fashion to the rolling performance we displayed earlier, we can look at the stability of a linear model's parameters through time. Figure 14 shows a three panel chart for the alpha, beta, and r-squared measures through time across a rolling window. Each chart calls **chart.RollingRegression** with a different **method** parameter.

Likewise, we can assess whether the correlation between two time series is constant through time. Figure 15 shows the rolling 12-month correlation between each of the peer group and the S&P500. To look at the relationships over time and take a snapshot of the statistical relevance of the measure, use **table.Correlation**, as shown in Figure 16.

3.9 Calculate Downside Risk

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 **Semi-Variance** would be a better measure of risk than variance. This measure is also called

Figure 14: Create a Rolling Regression

```
> charts.RollingRegression(managers[, c(manager.col, peers.cols),  
+   drop = FALSE], managers[, 8, drop = FALSE], rf = 0.03/12,  
+   colorset = redfocus, lwd = 2)
```

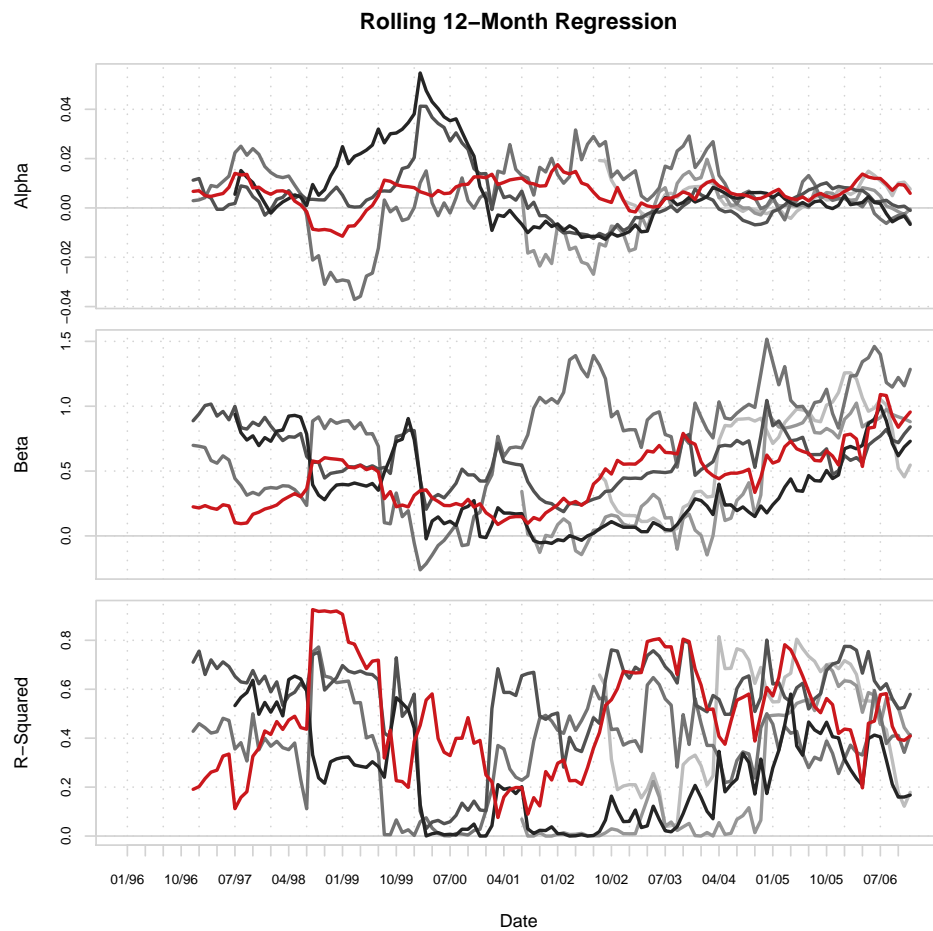


Figure 15: Chart the Rolling Correlation

```
> chart.RollingCorrelation(managers[, c(manager.col, peers.cols)],  
+   managers[, 8, drop = FALSE], colorset = tim8equal, lwd = 2,  
+   main = "12-Month Rolling Correlation")
```

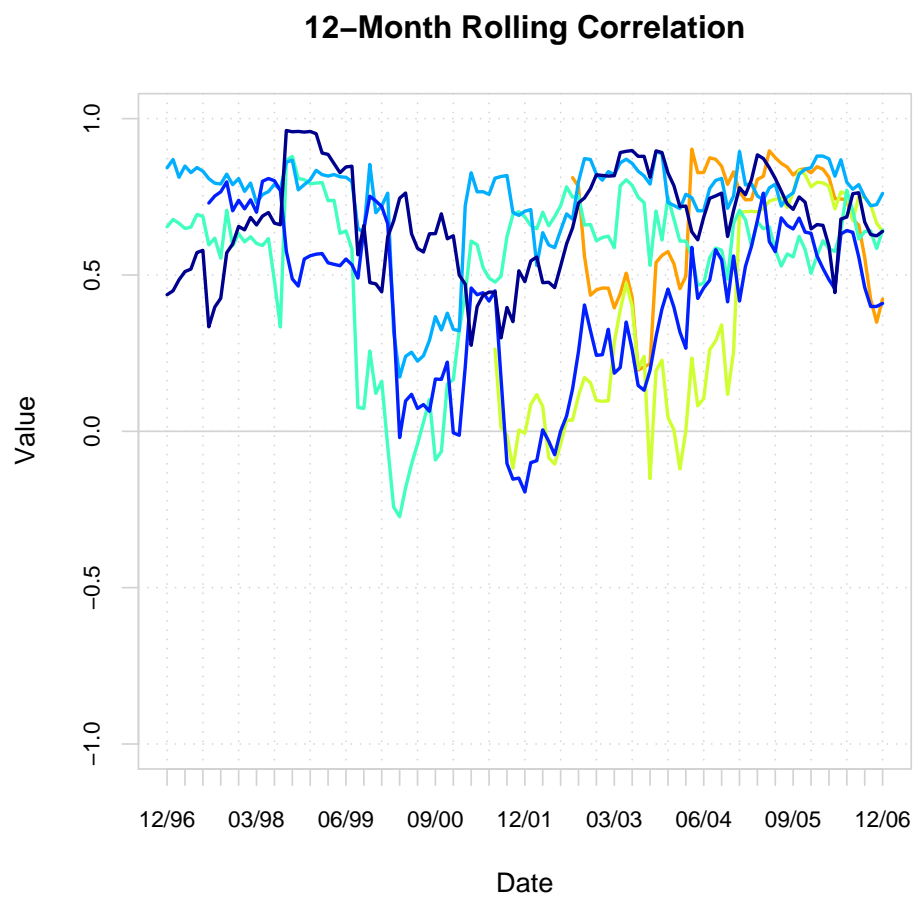


Figure 16: Calculate Correlations

```
> table.Correlation(managers[, c(manager.col, peers.cols)], managers[,
+      8, drop = F], legend.loc = "lowerleft")
```

	Correlation	p-value	Lower CI	Upper CI
HAM1 to SP500.TR	0.6688743	0.000000e+00	0.56226425	0.7536145
HAM2 to SP500.TR	0.4230678	8.862101e-07	0.26732059	0.5572733
HAM3 to SP500.TR	0.6617434	0.000000e+00	0.55345182	0.7480330
HAM4 to SP500.TR	0.5548691	5.065726e-12	0.42421745	0.6628915
HAM5 to SP500.TR	0.2806503	1.342386e-02	0.06047281	0.4747844
HAM6 to SP500.TR	0.5321481	6.023719e-06	0.32943057	0.6879665

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 `[quadprog]solve.QP` or `[tseries]portfolio.optim` or `[fPortfolio]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 use `table.Drawdowns` to find and sort them in order from worst/major to smallest/minor, as shown in Figure 18.

The `UpDownRatios` function may 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. Or, as mentioned above, you can also plot drawdowns with `chart.Drawdown`.

4 Conclusion

With that short overview of a few of the capabilities provided by `PerformanceAnalytics`, we hope that the accompanying package and documentation will partially fill a hole in the tools available to a financial engineer or analyst. If you think there's an important gap or possible improvement to be made, please don't hesitate to contact us.

Figure 17: Create a Table of Downside Statistics

```
> table.DownsideRisk(managers[, 1:6], rf = 0.03/12)
```

	HAM1	HAM2	HAM3	HAM4	HAM5	HAM6
Semi Deviation	0.0188	0.0203	0.0239	0.0397	0.0320	0.0173
Gain Deviation	0.0164	0.0347	0.0296	0.0314	0.0298	0.0157
Loss Deviation	0.0209	0.0099	0.0187	0.0371	0.0321	0.0132
Downside Deviation (MAR=10%)	0.0175	0.0168	0.0218	0.0386	0.0346	0.0152
Downside Deviation (rf=3%)	0.0151	0.0133	0.0188	0.0357	0.0316	0.0125
Downside Deviation (0%)	0.0142	0.0119	0.0176	0.0345	0.0303	0.0114
Maximum Drawdown	-0.1573	-0.2240	-0.2786	-0.2913	-0.3775	-0.0707
VaR (99%)	0.0696	0.0996	0.0985	0.1352	0.1075	0.0674
Beyond VaR	0.0704	0.1010	0.0997	0.1366	0.1078	0.0682
Modified VaR (99%)	0.0827	0.0804	0.0788	0.1282	0.0989	0.0523

Figure 18: Create a Table of Sorted Drawdowns

```
> table.Drawdowns(managers[, 1, drop = F])
```

	From	Trough	To	Depth	Length	To Trough	Recovery
1	Feb 2002	Sep 2002	Jul 2003	-0.15731054	18	8	10
2	May 1998	Aug 1998	Mar 1999	-0.13009220	11	4	7
3	Mar 2005	Apr 2005	Jul 2005	-0.03982848	5	2	3
4	Jun 1996	Jul 1996	Aug 1996	-0.03041080	3	2	1
5	Sep 2001	Oct 2001	Nov 2001	-0.02925960	3	2	1