

Portfolio Probe User's Manual

Burns Statistics

19th November 2010

Copyright 2003-2010 Burns Statistics Limited. All rights reserved.
<http://www.burns-stat.com/>
Edition 3: 2010 November 18

S-PLUS is a registered trademark of TIBCO.

Contents

1	Orientation	13
1.1	Overview of Functionality	13
1.2	Necessary Tools	13
1.3	Installing the Software	14
1.4	Loading the Software	14
1.5	Road Map	15
1.6	Typography Conventions	15
2	Random Portfolio Examples	19
2.1	Performance Measurement: Static	19
2.2	Testing the Effect of Constraint Bounds	21
2.3	Test a Risk Model	24
2.4	Performance Measurement: Shadowing	27
2.5	Evaluating a Trading Strategy	27
2.6	Going Farther	29
3	Generating Random Portfolios	33
3.1	The Command	33
3.2	Valuation	35
	Simple Use	35
	Collapsing Values	36
	Net Asset Value	37
	Random Weights	37
3.3	Working with Random Portfolios	38
	Small Selections	38
	Evaluating Portfolios	38
	Summary	39
3.4	Exporting Random Portfolios	40
	Writing monetary value	40
	Writing compact files	40
3.5	Combining Random Portfolio Objects	41
3.6	Unsatisfiable and Difficult Constraints	41
3.7	Adding a Utility Constraint	42
3.8	Going Farther	43

4	Constraints	45
4.1	Summary of All Constraints	45
	Round Lots	45
4.2	Monetary Value of the Portfolio	45
	Long-only Portfolios	47
	Long-short Portfolios	47
4.3	Limits on Assets	49
	max.weight	49
	universe.trade	50
	lower.trade and upper.trade	50
4.4	Number of Assets	51
	Number of Assets to Trade	51
	Number of Assets in the Portfolio	52
4.5	Threshold Constraints	52
	Trade Thresholds	52
	Portfolio Thresholds	53
	Summary of Threshold Inputs	53
4.6	Forced Trades	54
4.7	Positions	54
	Portfolio constraints	55
	Trade constraints	56
	Forced constraints	56
	Universe constraints	56
	Threshold constraints	56
	Tolerance	57
	Summary of positions inputs	57
4.8	Linear Constraints	57
	Building Constraints	58
	Bounds and lin.style	59
	Numerical Constraints: Risk Factors	61
	Numerical Constraints: Market Capitalization	62
	Mixing Numerical and Categorical Constraints	63
	Portfolio Constraints versus Trade Constraints	63
	Net Constraints versus Gross Constraints	63
	Long-side Constraints and Short-side Constraints	64
	Looking at the Effect of the Constraints	65
	Evaluating Un-imposed Constraints	66
	Inspecting Linear Constraints	66
4.9	Count Constraints	67
4.10	Alpha (Expected Return) Constraints	69
4.11	Variance Constraints	70
4.12	Tracking Error (Benchmark) Constraints	71
	Single Upper Bound	71
	Scaling	71
	Lower and Upper Bounds	71
	Multiple Benchmarks	72
	Advanced Use	72
4.13	Distance	72
	Alternative prices	73
	Multiple distances	73

4.14	Sums of Largest Weights	74
4.15	Cost Constraints	75
4.16	Number of Positions to Close	75
4.17	Quadratic Constraints	76
	Add Constraints to the Variance	76
	Impose Constraint Bounds	76
	Dummy Run	77
	Check for Benchmark	77
	Constraints out of Utility	78
	Actual Computation	78
4.18	Constraint Penalties and Soft Constraints	78
5	Details of Random Portfolio Examples	81
5.1	Performance Measurement: Static	83
5.2	Testing the Effect of Constraint Bounds	84
	Static Utility	84
	Caching Variances	85
	Dynamic Utility	85
5.3	Test a Risk Model	86
	Ex Ante Volatility Comparisons	87
	Ex Ante versus Realized Volatility	88
5.4	Performance Measurement: Shadowing	89
5.5	Evaluating a Trading Strategy	90
5.6	Going Farther	90
6	Optimizing Long-Only Portfolios	91
6.1	Required Inputs	91
	Monetary Value	91
	Utility	91
6.2	Examples for Passive Portfolios	92
	Minimize the Variance of the Portfolio	92
	Minimize Tracking Error	92
6.3	Examples for Active Portfolios	93
	Maximize the Information Ratio	93
	The Information Ratio with a Tracking Error Constraint	94
	Maximize Benchmark-relative Information Ratio	95
	Mean-Variance Optimization	96
	Mean-Volatility Optimization	96
	Buy-Hold-Sell List	96
6.4	Utility-free Optimization	97
6.5	Managing Cash Flow	98
	Injecting Money into a Portfolio	98
	Extracting Money out of a Portfolio	99
6.6	Asset Allocation	99
6.7	Going Farther	100

7	Optimizing Long-Short Portfolios	101
7.1	Required Inputs	101
	Monetary Value	101
	Utility	101
7.2	Examples	102
	Maximize the Information Ratio	102
	Maximize Return with a Bound on the Variance	103
	Minimize Variance Given a Long List and a Short List	103
	Mean-Variance Optimization	104
7.3	Managing Cash Flow	105
	Injecting Money into a Portfolio	105
	Extracting Money out of a Portfolio	105
7.4	Money Constraints	106
7.5	Real-Time Monitoring	106
7.6	Going Farther	108
8	General Use	109
8.1	Setting Up Data	109
	Prices and Other Imports	109
	Variance Matrix	110
	Adding a Benchmark to the Variance	112
8.2	The Random Generation or Optimization	112
8.3	Post-Optimization	113
	Explore the Trade	113
	Export the Trade	114
9	Trading Costs	115
9.1	Background	115
9.2	Specifying Costs	115
	Linear Costs	116
	Nonlinear Costs	116
	Polynomial Costs	117
9.3	Power Laws	118
9.4	On Scaling Costs Relative to Utility	119
9.5	Costs Due to Taxes	120
9.6	Going Farther	120
10	Practicalities and Troubleshooting	121
10.1	Easy Ways to Be Wrong	121
	Data Mangling	121
	Input Mangling	123
10.2	Suppressing Warning Messages	124
10.3	Cheatsheets	126
	Implied Ranges	126
	Threshold Inputs	127
	Positions Inputs	127
10.4	Troubleshooting	127
	Utility Problems	127
	Portfolio Problems	128
10.5	S Language Problems and Solutions	129

Creating Matrices	129
Debugging	129
11 Special Instructions	131
11.1 Special Instruction 1: Long-only when shorts exist	131
11.2 Special Instruction 2: Benchmark in long-short optimization	131
12 Adjusting Optimization Speed and Quality	133
12.1 Staying at a Given Solution	133
12.2 Reducing Time Use	134
12.3 The Optimization Process	135
12.4 Improving Quality	136
12.5 Testing Optimization Quality	136
13 Utility	139
13.1 Maximum Information Ratio	139
13.2 Mean-Variance Utility	140
13.3 Mean-Volatility Utility	141
13.4 Minimum Variance	141
13.5 Maximum Expected Return	141
13.6 Minimum Distance	142
13.7 Going Farther	142
14 Advanced Features	143
14.1 Multiplicity	143
14.2 Alpha and Variance Tables	144
14.3 Variance Constraints	145
14.4 Expected Return Constraints	147
14.5 Multiple Utilities	148
14.6 Utility Tables	149
14.7 Multiplicity Examples	150
Dual Benchmarks	150
Benchmark-relative Utility and Absolute Variance Constraint	152
Rival Variance Forecasts	153
Multiple Time Periods	154
Credit Risk	155
Multiple Scenarios	156
14.8 Compact Variance Objects	161
The Variance List	161
15 Dregs	163
15.1 Portfolio Probe Constituents	163
15.2 The Objectives	163
15.3 Writing C or C++ Code	163
15.4 Bug Reporting	164

List of Tables

2.1	Simple analyses of the MACD funds.	20
2.2	Performance measurement using initial positions.	28
4.1	Summary of constraints.	46
10.1	Implied ranges of arguments.	127
10.2	The meaning of threshold inputs.	127
10.3	The column order for the positions argument.	128
14.1	Arguments for multiple variances and expected returns.	144
14.2	Utilities codes for the utility table.	150

List of Figures

1.1	Some suggested routes through the document.	16
1.2	Possible route for utility-free optimization.	17
2.1	MACD funds compared to random portfolios: 2007.	20
2.2	MACD funds compared to random portfolios: 2007 and 2008.	21
2.3	Mean-variance utility for calendar year 2007.	22
2.4	Approximate minimum volatility.	23
2.5	Mean-variance utility from overlapping 60-day windows.	23
2.6	Shrinkage versus factor model estimates of volatility for 2007 Q4.	25
2.7	Shrinkage versus factor model estimates of volatility for 2008 Q4.	25
2.8	Quarterly correlations shrinkage versus realized: long-only.	26
2.9	Quarterly correlations shrinkage versus realized: dollar neutral.	26
2.10	Quarterly correlations shrinkage versus realized: 120/20.	27
2.11	Performance measurement: 200% turnover for 2007.	28
2.12	Performance measurement: 200% turnover for 2007-2008.	28
2.13	MACD efficacy for Q1 of 2008.	30
2.14	MACD efficacy for the first half of 2008.	30
2.15	MACD efficacy for the year 2008.	31
2.16	Outperformance of MACD for twenty initial portfolios.	31
2.17	Mean outperformance of MACD over twenty initial portfolios.	32
7.1	Constraints on gross, net, long and short values.	107

Chapter 1

Orientation

This chapter has diverse aims:

- It provides a brief overview of Portfolio Probe functionality.
- It explains what software you need in order to run Portfolio Probe.
- It suggests a route through the rest of this document, given a task and a state of mind.
- It presents the typographic conventions of the document.

1.1 Overview of Functionality

The two primary aims of Portfolio Probe are:

- To generate random portfolios (the `random.portfolio` function).
- To optimize a portfolio (the `trade.optimizer` function).

The two functions have the same inputs (except for saying how many random portfolios you would like, and whether you want random portfolios or random trades).

All of the rest of Portfolio Probe is support for these two tasks.

1.2 Necessary Tools

You need to choose a language in which to run Portfolio Probe. It can be one of three:

- R, which can be downloaded for free via:
<http://www.r-project.org/>

There are some commercial distributions of R as well.

- S-PLUS, sold by TIBCO:
<http://spotfire.tibco.com/>

- C or C++. You can call Portfolio Probe functionality in a program that you write.

The last option of using C is not recommended—it requires considerable effort, and likely has little or no benefit.

S-PLUS and R are versions of the S language, and Portfolio Probe has been written to work with either version of S. This document assumes you are using S (as opposed to using C code). Portfolio Probe uses only very general features of S so it should run the same in any version of R or S-PLUS.

When this document says “S”, it means either R or S-PLUS—the term “S” should not be construed to mean only S-PLUS. Some of the examples explicitly assume R is being used—the same effect would be done slightly differently in S-PLUS.

Programming experience is not mandatory—whatever your objective, there is likely to be an example provided in this manual that is close to your case.

The present document assumes knowledge of S to the level of “Some Hints for the R Beginner”—a brief, structured introduction which can be found in the Tutorials section of <http://www.burns-stat.com/>. Commands beyond that level are included and explained.

While it is reasonably easy to start using Portfolio Probe, there is a lot of room to grow. Portfolio Probe’s flexibility and the power of S can carry you a long way.

1.3 Installing the Software

If you are using R and your machine has access to the internet, then you can install Portfolio Probe with:

```
> install.packages("PortfolioProbe",  
+   repos="http://www.portfolioprobe.com/R")
```

See the Frequently Asked Questions in the “User area” of www.portfolioprobe.com if a firewall blocks this command or if you are using S-PLUS.

1.4 Loading the Software

If Portfolio Probe was installed in the default place, then Portfolio Probe is loaded into an R session with:

```
> library(PortfolioProbe)
```

It is possible that the `lib.loc` argument to `library` may be required. In S-PLUS it will depend on the particular installation, but something similar is likely.

1.5 Road Map

We suggest that there is one of two frames of mind that you are likely to have:

- Conceptual: primarily wanting to understand the task
- Operational: primarily wanting to do the task

With the choice of two tasks, that produces four possible routes through the document. Figure 1.1 is a graphical view of the suggestions.

Note that not all chapters appear in the maps. A lot of the document can be reserved for reference as the need arises. In particular if you are only generating random portfolios, you can safely ignore several of the chapters.

If you are only interested in utility-free optimization, then Figure 1.2 shows a possible route through this document. This supposes that you will impose a turnover constraint rather than providing trading costs for the assets.

1.6 Typography Conventions

Computer commands or pieces of commands are written in **this font**. For example, a variance argument is written as **variance** whenever it is the argument itself that is referred to. Entire commands are written in the same font. Commands are introduced by “>” (which is the prompt in the S language) so that any output can be distinguished from the input. An example is:

```
> rep(0, 6)
[1] 0 0 0 0 0 0
```

The user types “rep(0, 6)” (followed by return), and the next line is the response from S.

Commands may span more than a single line—the second and subsequent lines of a command are introduced by “+”. For example:

```
> op <- trade.optimizer(prices, varian, gross.value=1e6,
+   long.only=TRUE)
```

The second line of the command starts with `long.only` (the “+” is not typed by the user, but rather by S). There is no output in this example.

S code note

The only catch with multi-line commands is that it needs to be clear to S that the command is incomplete. In this example the command needs a closing parenthesis.

Occasionally a fragment of code is written, in which case there are no introductory prompts.

In addition to S code notes, there are boxes which contain “cautions” and “notes”.

Figure 1.1: Some suggested routes through the document.

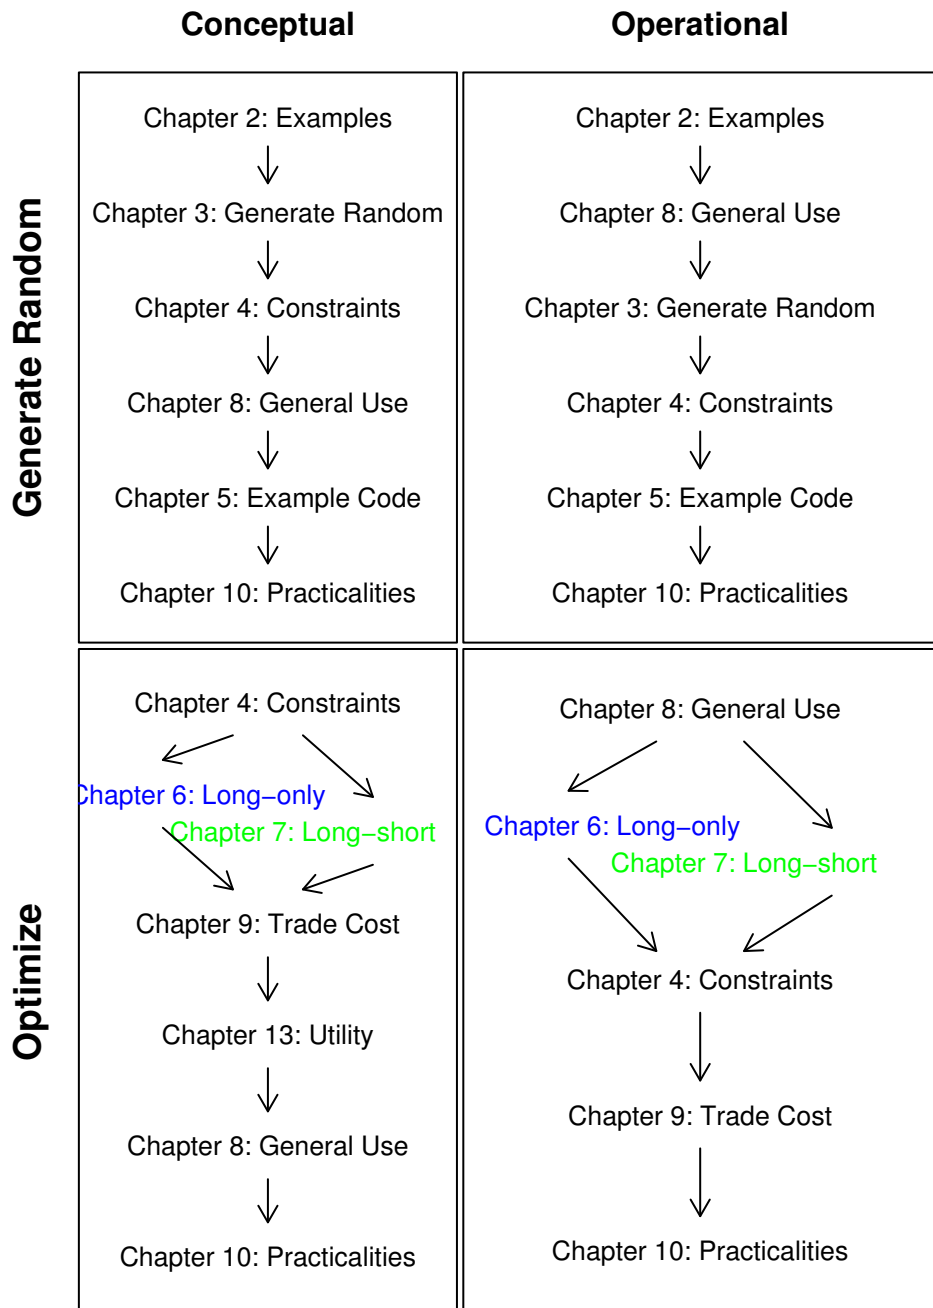
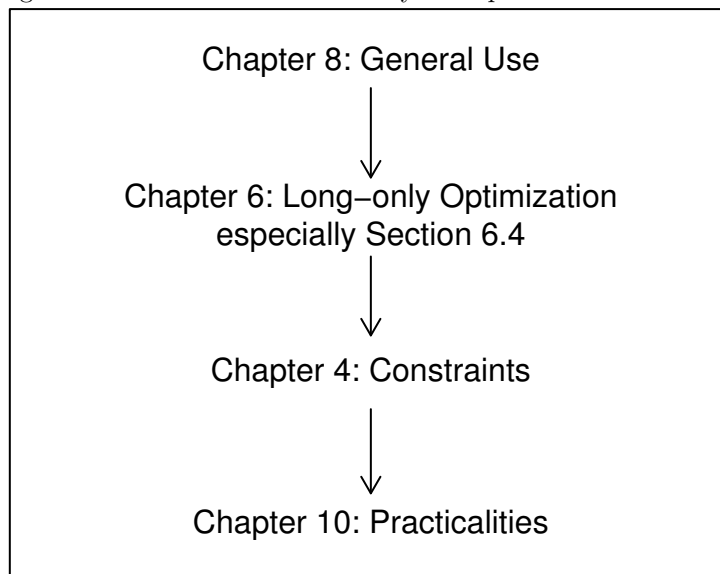


Figure 1.2: Possible route for utility-free optimization.



Chapter 2

Random Portfolio Examples

This chapter shows some of the more common applications of random portfolios. Chapter 5 contains the code that created each of these examples.

These examples use daily (adjusted) price data for an arbitrary but non-random selection of 70 US large cap stocks for the years 2006 through 2008.

2.1 Performance Measurement: Static

Suppose we have a particular fund that we would like to evaluate. We can generate a number of random portfolios that obey the constraints of the fund. The fund manager *might* have selected any one of the random portfolios that we generate. We can compare the selection the fund manager did make with the selections of the proverbial monkeys.

The static method of performance measurement is to create a set of random portfolios that satisfy the fund's constraints at the start of the period, assume that we hold these portfolios throughout the time period, and compare the return of the actual portfolio to the distribution of returns from the random portfolios.

The information that we need in order to apply this method is:

- the constraints under which the fund operates

While it would be best if we knew the constraints exactly, approximations of the constraints can be used.

For the examples of this section two “funds” were created by starting with a specific portfolio at the start of 2007 and using the constraints:

- long-only
- the number of assets in the portfolio can be no more than 20
- the maximum weight for any asset is 10%

Every day throughout 2007 and 2008 an optimal trade was made according to an MACD signal (the default from the `MACD` function in the `TTR` R package

Figure 2.1: MACD funds compared to random portfolios: 2007.




Table 2.1: Simple analyses of the MACD funds.

	200% turnover	400% turnover
2007 return	25.8%	28.6%
2007 performance percentile	9	4
2007-2008 return	-32.7%	-24.2%
2007-2008 performance percentile	97	69

[Ulrich, 2009]). Note that this is meant neither to praise nor to pick on MACD—it is merely an easy way of producing expected returns that you can replicate yourself if you wish. The difference between the two funds is that one has a target of 200% annual turnover (buys plus sells) while the other has a 400% turnover target.

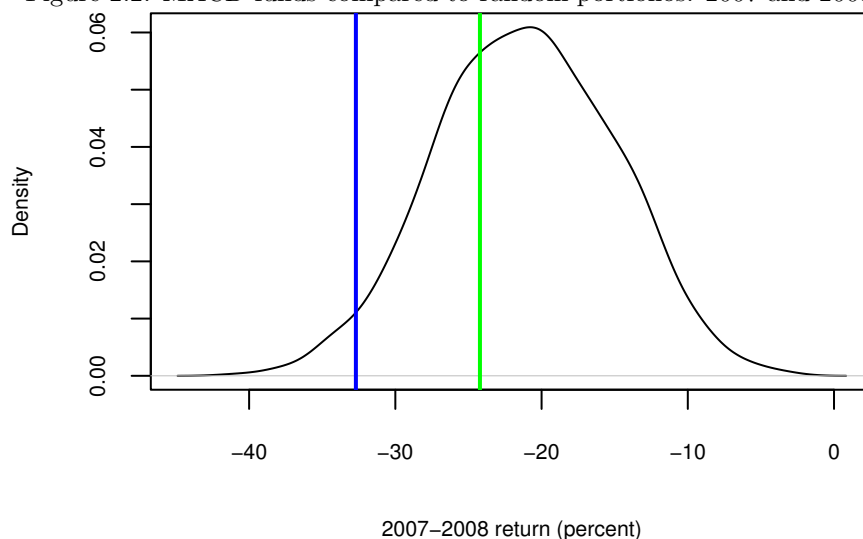
Figure 2.1 shows the results for 2007. Both funds did reasonably well. The situation changes if we extend the analysis through 2008—Figure 2.2 displays those results.

Table 2.1 shows the numerical values for the analyses. The percentile shows where the return of the actual fund falls among the returns of the random portfolios. By convention, performance percentiles are near zero for good results and near 100 for bad results.

One failing of this analysis is that the random portfolios are not guaranteed to obey the constraints at the end of the period. Even so, the results are likely to be useful, and the effort to create the analysis is minimal. The more elaborate analysis in Section 2.4 overcomes the constraint violation problem as well as being substantially more sensitive.

The number of random portfolios that were generated was 1000. This is not

Figure 2.2: MACD funds compared to random portfolios: 2007 and 2008.



enough to make the density plots completely smooth, but is more than enough for the substantive part of the task at hand.

2.2 Testing the Effect of Constraint Bounds

Constraints are put on portfolios as a form of insurance. We are willing to give up some up-side in order to be protected from down-side. In general, we are striking that bargain blindly—we don't know the up-side we are forgoing, and we don't know the extent of the protection we are getting. Random portfolios allow us to examine the trade-off over historical periods.

For the examples of this section, the constraints that are imposed are:

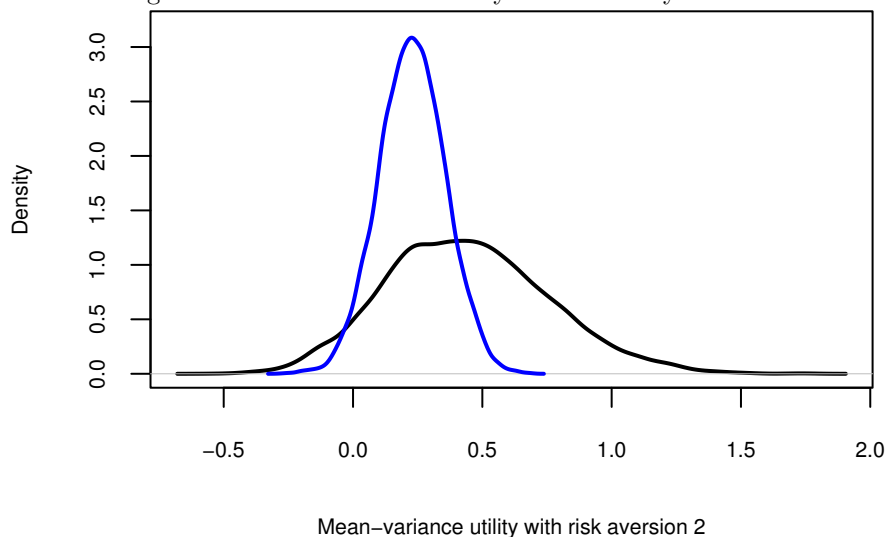
- long-only
- no more than 30 assets
- no asset with a weight greater than 10%

and some times

- volatility no more than some value

Figure 2.3 shows the distribution of mean-variance utility with risk aversion 2 for portfolios that are held throughout 2007. The blue line indicates the distribution for portfolios that have the additional constraint that the volatility is no more than 8% (as measured *ex ante* at the beginning of the year). We see that we lose a lot of up-side with the volatility constraint. But this is merely a snapshot—it tells us what happened to have occurred, but it doesn't indicate what we can expect to occur.

Figure 2.3: Mean-variance utility for calendar year 2007.

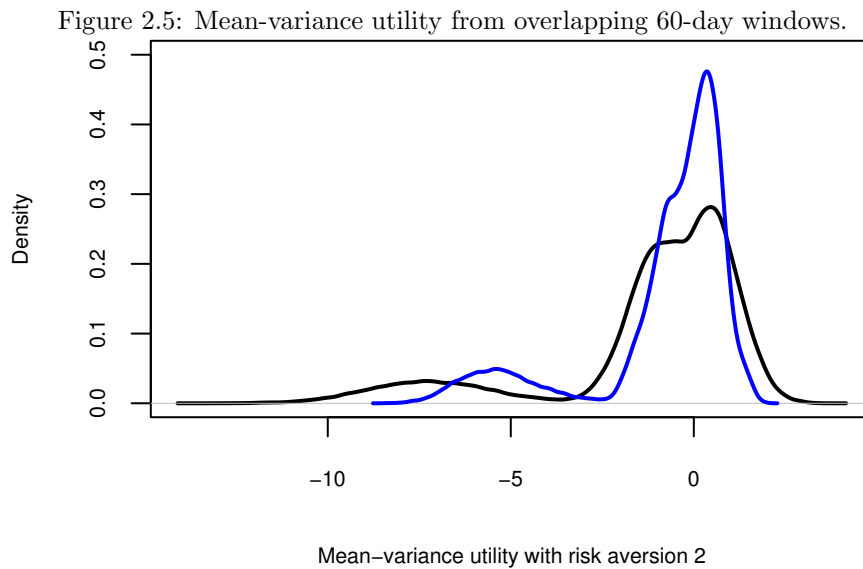



Suppose that we are interested in utility over a quarter (or, for convenience, 60 trading days). We can move that 60-day window one day at a time and get different answers. We have two years of data (2007 and 2008) in our example, so we have roughly 440 windows. There are at least two approaches. One is to form the portfolios with their constraints at the start of 2007 and never change them. That is easy to do, but probably doesn't answer the question that we really have in mind.

The second approach forms random portfolios at the start of each window and uses these to calculate the utility that is realized for the window. Often the constraints will be static throughout, even though the portfolios change with each window. In this case, imposing a static volatility constraint is a problem—we start by looking at what the 8% volatility constraint looks like in the broader context.

Figure 2.4 shows the minimum (*ex ante*) volatility that is achievable throughout the time period we are using. Clearly we can't always impose a constraint that volatility is no more than 8% since the minimum possible is greater than 8% for most of the period. The 8% constraint on the first day is approximately 120% of the minimum volatility. So we'll use 120% of the minimum volatility as the upper limit on volatility when it is constrained.

Figure 2.5 compares the rolling distributions of utility with and without the volatility constraint. We see bimodal distributions. There is what might be considered a normal mode and a disaster mode. In normal mode, the volatility constraint reduces the up-side about the same amount as the down-side. In disaster mode, the volatility constraint is very useful. So this looks quite good for the volatility constraint. Note, though, that this is quite a short time period from which to make generalizations.



2.3 Test a Risk Model

We have a wide choice of risk models. If we could compare results of risk models on a large number of realistic portfolios, we could make informed decisions about the best risk models for particular tasks. Random portfolios can provide those realistic portfolios in a convenient form. This section merely provides a brief taster of what is possible.

The `BurStFin` R package contains two functions for estimating variance matrices (see page 110). One estimates a statistical factor model, the other shrinks the sample variance towards the equal correlation matrix. Each function was given 250 days of returns and used default values of all arguments. Since there are only 70 assets, this is a rather atypical case for these estimators. Often there are more assets than time points. Even so, the amount of shrinkage towards equal correlation ranged from about 17% to about 32%.

One thousand random portfolios were generated at the start of each of eight quarters. The constraints were:

- long-only
- no more than 30 assets in the portfolio
- 5% maximum weight for each asset
- the sum of the 4 largest weights is no more than 15%

Throughout this section the last three constraints are maintained, but the first is sometimes changed.

Figure 2.6 compares the *ex ante* estimates of volatility from the two models with data up to the start of the fourth quarter of 2007. The estimates agree very well. Figure 2.7 is a similar plot, but for the fourth quarter of 2008.

We would like to see high correlations between the estimated and realized volatilities for the portfolios. Figure 2.8 shows the correlation for each quarter between the shrinkage estimate of volatility and the realized volatility. The red dashed lines represent 95% confidence intervals for the correlation as found by a statistical bootstrap. The pictures are very similar for factor models.

As we need to expect, the predicted level of volatility in the last quarter of 2008 is severely lower than the realized volatility—the actual volatilities tended to be more than twice the predicted.

Two more sets of random portfolios were generated. One set was long-short dollar neutral (zero net value). The other was 120/20 long-short. In both cases the same constraints of 30 assets, 5% maximum weight and 15% sum of the 4 largest weights were imposed. Figure 2.9 shows the quarterly correlations between the *ex ante* shrinkage estimate and the realized volatility for the dollar neutral portfolios. Figure 2.10 shows the correlations for the 120/20 portfolios.

Random portfolios give you the tools to select the best variance estimators for the types of portfolios that you hold, and allow you to see how good that is.

Figure 2.6: Shrinkage versus factor model estimates of volatility for 2007 Q4.




Figure 2.7: Shrinkage versus factor model estimates of volatility for 2008 Q4.

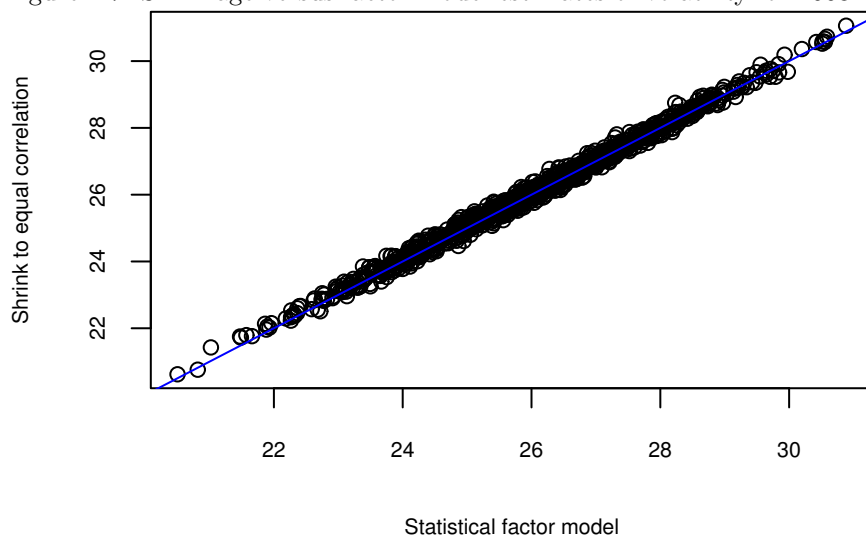


Figure 2.8: Quarterly correlations shrinkage versus realized: long-only.




Figure 2.9: Quarterly correlations shrinkage versus realized: dollar neutral.





Figure 2.10: Quarterly correlations shrinkage versus realized: 120/20.



2.4 Performance Measurement: Shadowing

Section 2.1 presents the static method of performance measurement. While that method is more powerful than peer groups and benchmarks, we can often do even better. If we know some or all of the positions of the portfolio at the beginning of the test period, random portfolios can give us a much more sensitive test of performance.

The information that we need for this method is:

- the constraints under which the fund operates
- some or all of the positions at the start of the time period
- the known or approximate turnover of the fund throughout the period

With this method we imitate the trading of the fund throughout the period. We get one random path by means of a series of random trades starting with what we know about the initial portfolio. We then repeat that a number of times.

Let's revisit the example we saw in Section 2.1. We know the full portfolio at the start of the period, and we know the typical daily trading. We use this knowledge to get the distribution of returns with zero skill.

Figure 2.11 shows the results for the year 2007 with 200% turnover, while Figure 2.12 shows the 2007-2008 results for 200% turnover. The strategy is very clearly outperforming chance. The results are even more extreme for 400% turnover. Table 2.2 shows the numerical results.

2.5 Evaluating a Trading Strategy

It is quite common to perform a backtest of a strategy. However, the backtest itself only shows how much the strategy makes or loses in the given situation. It

Figure 2.11: Performance measurement: 200% turnover for 2007.




Figure 2.12: Performance measurement: 200% turnover for 2007-2008.

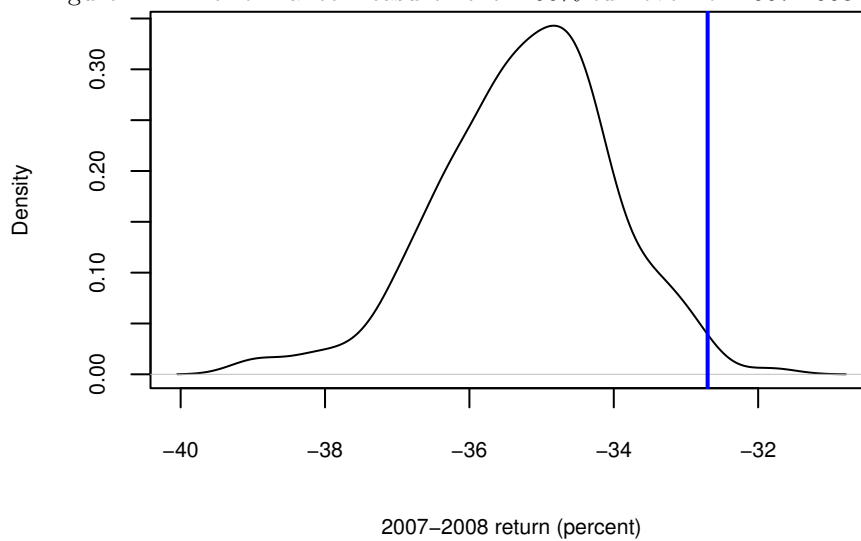


Table 2.2: Performance measurement using initial positions.

	200% turnover	400% turnover
2007 return	25.8%	28.6%
2007 performance percentile	0	0
2007-2008 return	-32.7%	-24.2%
2007-2008 performance percentile	1	0

doesn't provide any information on whether it is luck or skill. Random portfolios can be used to help decide that.

The previous section discussed performance measurement using the shadowing method. We can do the same thing to test a trading strategy. One key difference is that we are doing it *ex ante* rather than *ex post*. Plus, we know the constraints and the trading exactly. Another difference is that with performance measurement we are interested in the particular portfolio that was actually held at the start of the period, but with evaluating a strategy we want the performance to be good no matter where we start.

Operationally the real difference is that when evaluating a trading strategy we create a set of random portfolios at which to start and we see how the strategy does during the time period starting at each of those portfolios. That is, we repeat the shadowing process a number of times.

In the example we create 20 initial portfolios at the beginning of 2008, and see how the MACD strategy performs after 1 quarter, 2 quarters and the whole year. We have the constraints:

- long-only
- no more than 20 names in the portfolio
- no asset weight greater than 10%
- 200% turnover target


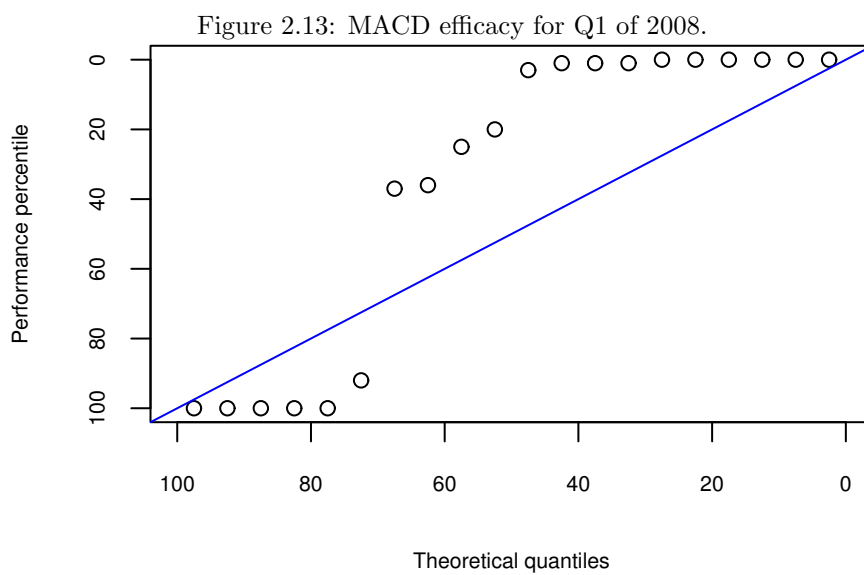
Figure 2.13 shows how good MACD was for these constraints during the first quarter of 2008. It has close to zero value on average—that is, it is neither good nor bad. However, for a specific starting portfolio it tends to either significantly outperform or significantly underperform. If the strategy truly had no effect, then the percentiles would fall roughly along the diagonal line. That the percentiles tend to be zero or one hundred is an indication of the high power of the technique. (Each strategy is compared to only 100 random runs, but adding more random runs would at most move the percentiles slightly.)

Figure 2.14 shows how the strategy did for the first half of the year. It did exceptionally well—the strategy outperformed all random paths for all but one of the starting portfolios; in the exception it beat almost all the random paths. The picture for the whole year (Figure 2.15) is almost at the other extreme. For the whole year the strategy significantly underperformed with almost all of the starting points.

From those plots we don't know if the amount of outperformance or underperformance is large or trivial. To learn that, we can look at the return of the strategy relative to the mean return of the random paths. That is, we compare the strategy return to what is expected if there were zero skill. Figure 2.16 shows the outperformance for each of the twenty runs. Figure 2.17 plots the mean outperformance over the twenty runs.

2.6 Going Farther

- Chapter 3 discusses the function that generates random portfolios.



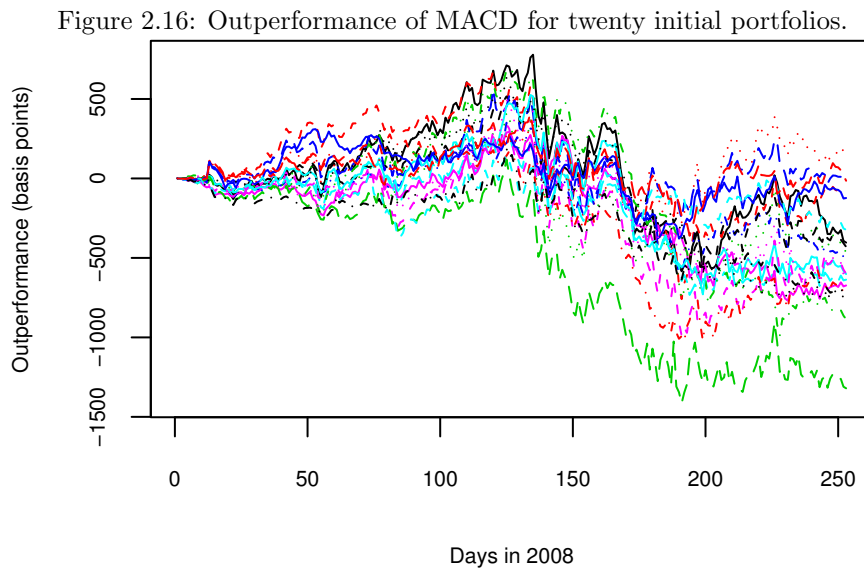




Figure 2.17: Mean outperformance of MACD over twenty initial portfolios.



- Chapter 4 discusses the constraints that may be imposed.
- Chapter 5 shows the commands that created the examples that were displayed in this chapter.

Chapter 3

Generating Random Portfolios

The `random.portfolio` function generates a list of portfolios (or the trades for them) that satisfy constraints but pay no attention to utility.

3.1 The Command

To generate random portfolios you give `random.portfolio` the number of random portfolios that you want to generate, the basic information for the problem, and the constraints that you would like. There is also the `out.trade` argument which controls whether it is the random portfolio (the default) or the trade which is output.

At a minimum you need to specify the vector of prices and the amount of money in the portfolio. One possibility is:

```
> randport1 <- random.portfolio(prices=prices, long.only=TRUE,
+   gross.value=1e6)
```

Of course this is not very interesting. You are likely to want more than one random portfolio, and to have non-trivial constraints imposed.

S code note

The notation `1e6` means 1,000,000, that is, a one followed by six zeros.

The prices (always required) needs to be a vector of positive numbers that has names which identify the assets in the problem. Here is an example of the first few values of a suitable price vector:

```
> head(pricevec)
stockA stockB stockC stockD stockE stockF
27.63 19.46 11.67 5.79 5.15 20.99
```

The assets named in `prices` define the universe of assets for the problem. In examples the `prices` argument is usually given a vector called `prices`—in reality the name of the vector can be whatever you like.

The other two pieces of “basic information” are the variance matrix and the vector of expected returns—neither of these are required. You only need to give these if they are involved in a constraint.

To generate 100 random portfolios that have country and sector constraints, no more than a 4% tracking error and no more than 55 assets, the following command would do:

```
> randport2 <- random.portfolio(100, prices, varian,
+   long.only=TRUE, gross.value=1e6,
+   bench.constraint = c(spx=.04^2/252),
+   port.size=55, lin.constraints=cntrysect.constraint,
+   lin.bounds=cntrysect.bounds)
```

S code note

The first three arguments in the call that creates `randport2` do not need to have the name of the argument specified because they are all in the order of the arguments in the definition of the function. In contrast the call that creates `randport1` uses the argument name in all cases. If there is any doubt, then it is safest to give the argument by name.

The examples so far assume that there is no existing portfolio (or that it doesn't matter). The `existing` argument gives the current portfolio.

```
> randport3 <- random.portfolio(100, prices, varian,
+   long.only=TRUE, gross.value=1e6,
+   bench.constraint = c(spx=.04^2/252),
+   existing=current.portfolio,
+   port.size=55, lin.constraint=cntrysect.constraint,
+   lin.bounds=cntrysect.bounds)
```

Sometimes it is more convenient to have the trades rather than the portfolios. If you want the trades, just set the `out.trade` argument to `TRUE`:

```
> randtrade3 <- random.portfolio(100, prices, varian,
+   long.only=TRUE, gross.value=1e6,
+   bench.constraint = c(spx=.04^2/252),
+   existing=current.portfolio,
+   port.size=55, lin.constraint=cntrysect.constraint,
+   lin.bounds=cntrysect.bounds, out.trade=TRUE)
```

S code note

The full name of the `out.trade` argument must be given. This is unlike almost all other arguments where only enough of the first portion of the name needs to be given to make it unique among the arguments to the function. (For

a full explanation of argument matching in S, see [Burns, 1998] page 19 or [Burns, 2009].)

If the `existing` argument is not given or is `NULL`, then it doesn't matter which value `out.trade` has—the output is the same in either case.

The result of a call to `random.portfolio` is a list where each component of the list is a portfolio (or trade). The object has a number of attributes including a `class` attribute ("`randportBurSt`"). Here is a small example:

```
> random.portfolio(2, priceten, gross.value=1e5,
+   long.only=TRUE, port.size=3, max.weight=.5)
[[1]]
stockA stockB stockJ
  1337   1481   6484
[[2]]
stockB stockF stockH
  1274   2382   3100
attr(,"call")
random.portfolio(number.rand = 2, prices = priceten,
  gross.value = 1e+05, long.only = TRUE, port.size = 3,
  max.weight = 0.5)
attr(,"timestamp")
[1] "Thu Oct 29 11:36:03 2009"
attr(,"class")
[1] "randportBurSt"
seed attribute begins: 1 -1142929704 1716596987 -285978235
```

Each component of the list is a portfolio (or trade), which is a vector giving the number of asset units (shares, lots, contracts) for each asset that appears.

3.2 Valuation

Random portfolios are seldom interesting in their own right, generally it is necessary to manipulate them to get the information we want. This section and the next discuss several common tasks.

Simple Use

The `valuation` method for random portfolios can give a variety of results relating to the monetary content of the random portfolios. The simplest usage is to give a vector of prices in addition to the random portfolio object:

```
> head( valuation(ranport, pricevec) , 2)
[[1]]
  stockB   stockC   stockJ
   38.92 50461.08 -49500.00
[[2]]
  stockB   stockI   stockJ
-28625.66 50497.42 -20871.84
```

What we get is the amount of money in each position in each portfolio. Any vector of prices (that contains the assets in the random portfolio object) can be used. It can be the price vector used to create the random portfolios, or some other price vector.

A matrix of prices can be given as well. The columns should correspond to assets:

```
> head( valuation(ranport, pricemat) , 2)
[[1]]
      stockB  stockC stockJ
20091012  38.92 50461.08 -49500
20091013  38.34 52882.52 -46875
20091014  36.20 53574.36 -48750
[[2]]
      stockB  stockI  stockJ
20091012 -28625.66 50497.42 -20871.84
20091013 -28199.07 50358.50 -19765.00
20091014 -26625.10 52650.68 -20555.60
```

The price matrix we used looks like:

```
> pricemat[, 1:4]
      stockA stockB stockC stockD
20091012  27.63  19.46  11.67   5.79
20091013  29.47  19.17  12.23   5.85
20091014  27.86  18.10  12.39   5.43
```

Collapsing Values

Often we don't care about the individual positions, but we do care about the gross value of the portfolios. Use the `collapse` argument to get this:

```
> valuation(ranport, pricevec, collapse=TRUE)
[1] 100000.00 99994.92 99992.25 99993.29
```

This has length 4 because that is the number of portfolios in `ranport`.

Perhaps the most useful is to use `collapse` and a matrix of prices:

```
> valuation(ranport, pricemat, collapse=TRUE)
      [,1] [,2] [,3] [,4]
20091012 100000.00 99994.92 99992.25 99993.29
20091013 99795.86 98322.57 101966.21 106510.98
20091014 102360.56 99831.38 104858.49 107995.59
```

The result is a matrix with as many rows as the price matrix and as many columns as there are portfolios in the random portfolio object.

It is possible to get the net value rather than the gross:

```
> valuation(ranport, pricemat, collapse=TRUE, type="net")
      [,1] [,2] [,3] [,4]
```

```

20091012 1000.00 999.92 999.93 999.99
20091013 6045.86 2394.43 3872.21 -2285.54
20091014 4860.56 5469.98 2364.57 -1377.59

```

The `type` argument is only used when `collapse` is `TRUE`.

Net Asset Value

The value of a long-only portfolio is generally unambiguous. That is not true of long-short portfolios. For long-short portfolios there is not only the issue of the value within the portfolio, but also how much cash is supporting the portfolio. This matters when valuing the portfolio in order to compute returns and realized volatility.

The `type="nav"` and `cash` arguments allow control over this. The `cash` argument (which is only consulted when the `type` is `"nav"`) can either be a single number or a vector with length equal to the number of rows in `prices`. When `type` is `"nav"` and `cash` is not given (or is `NULL`), then it is computed to be the gross value at the first time point—this is valid for some long-short portfolios but certainly not all.

Here are some plausible calls:

```

> valuation(randport.dollar.neutral, pricemat, collapse=TRUE,
+   type='nav') # no leverage
> valuation(randport.dollar.neutral, pricemat, collapse=TRUE,
+   type='nav', cash = some.value)
> valuation(randport.120.20, pricemat, collapse=TRUE,
+   type='nav', cash = 0)

```

Random Weights

If you want to have the random portfolios expressed in terms of weights rather than asset units, then you can use the `valuation` function with the `weight` argument set to `TRUE`. For example:

```

> head( valuation(ranport, pricemat, weight=TRUE) , 2)
[[1]]
      stockB  stockC  stockJ
20091012 0.0003892000 0.5046108 -0.4950000
20091013 0.0003841843 0.5299070 -0.4697089
20091014 0.0003536518 0.5233887 -0.4762577
[[2]]
      stockB  stockI  stockJ
20091012 -0.2862711 0.5049999 -0.2087290
20091013 -0.2868016 0.5121764 -0.2010220
20091014 -0.2667007 0.5273961 -0.2059032

```

This works whether `prices` is just a vector or a matrix.

3.3 Working with Random Portfolios

Small Selections

You can use `head` to get the first few random portfolios, and `tail` to get the last few. These are generic functions in R. Their random portfolio methods return an object that retains the class and other attributes of the original object.

These functions can be useful to inspect the portfolios to see if they look reasonable without printing hundreds or thousands of portfolios to the screen. They can also be used to test commands, such as the example immediately below.

Evaluating Portfolios

The sister function to `random.portfolio` is `trade.optimizer`. It can be of interest to see some of the values that the optimizer would return for each of the random portfolios. The `randport.eval` function does that: for each of the random portfolios (or trades) in the object it finds what the optimizer says about it. You can select which components of the output of `trade.optimizer` to keep (using the `keep` argument). The result is a list as long as the random portfolio object and each component of that list is a list containing the kept components.

Here is a small example of keeping the portfolio variances:

```
> randport.eval(head(randport4, 3), keep='var.values')
[[1]]
[[1]]$var.values
[1] 382.3576
[[2]]
[[2]]$var.values
[1] 147.6476
[[3]]
[[3]]$var.values
[1] 134.6368
```

In this case where we are returning only one number per portfolio, it makes more sense to coerce this to a numeric vector:

```
> unlist(randport.eval(head(randport4, 3), keep='var.values'),
+       use.names=FALSE)
[1] 382.3576 147.6476 134.6368
```

Keep in mind that these values are *ex ante* predictions—they may or may not have much relation to realized variance.

note

In `randport.eval` the optimizer is called using the same names of objects as was used when the random portfolio object was originally created. Objects with these names must be visible at the time that `randport.eval` is used. If any of

these objects has changed, then it is the current value rather than the original value that is used.

caution

Additional arguments or changes to arguments may be given to `randport.eval` so that what the optimizer does is not exactly what `random.portfolio` did. If you are making a change to an argument, then you need to use the exact same abbreviation (if any) as in the original call to `random.portfolio`.

There is a `FUN` argument to `randport.eval` that, if given, applies that function to each of the portfolio objects that are created. For example, we could do:

```
> randport.eval(head(randport4, 3), FUN=summary)
```

Or perhaps a more useful command along the same lines:

```
> do.call("rbind", randport.eval(randport4,
+   FUN=function(x) summary(x)$number.of.assets))
```

S code note

The command:

```
> do.call("rbind", some.list)
```

is equivalent to the command:

```
> rbind(some.list[[1]], some.list[[2]], ...,
+   some.list[[length(some.list)]])
```

Summary

The `summary` method for random portfolios shows how many assets are in the portfolios, and the number of times each asset appears in a portfolio:

```
> summary(randport5)
$port.size
 7  8  9 10
 1 32 165 802
$count.assets
stockC stockD stockA stockB stockE stockF stockI
 1000  1000   987   975   973   972   972
stockG stockH stockJ
  968   961   960
```

This shows us that out of the 1000 portfolios, 802 contained all 10 assets, 165 had 9 assets, 32 had 8 assets and 1 had 7 assets. We also see that `stockC` and `stockD` were both in all of the portfolios while `stockJ` was only in 960 of them.

3.4 Exporting Random Portfolios

The `deport` function will write files containing the result of `random.portfolio`. The simplest use is:

```
> deport(randport2)
[1] "randport2.csv"
```

This writes a comma-separated file where the columns each correspond to one of the assets that appear in the object and the rows correspond to the portfolios or trades. There are arguments that allow you to switch the meaning of rows and columns, and to give a universe of assets (which must include all of those appearing in the object). See the help file for details.

Writing monetary value

If you want the file to represent money rather than the number of asset units, you can use the `multiplier` argument:

```
> deport(randport2, multiplier=prices, file="randval1")
[1] "randval1.csv"
```

Writing compact files

If you would like a compact form of the file when there is a large universe of assets, you can use the `append` argument and iteratively write each portfolio to the file.

```
> deport.randportBurSt(randport2[1], file="randcompact")
> for(i in 2:length(randport2)) {
+   deport.randportBurSt(randport2[i], file="randcompact",
+     append=TRUE)
+ }
```

We start by first writing just the first portfolio with the default value of `append`, which is `FALSE`. The full name of the `deport` function needs to be used because the result of the subscripting has lost the `class` attribute. Specifying the file name is a necessity in this case. The remaining portfolios are written by a loop with `append` set to `TRUE`. Note that the loop is assuming that there are at least 2 random portfolios in `randport2`.

S code note

If the switch from `deport` to `deport.randportBurSt` confuses you, here is an explanation. The `deport` function is magic in that it is a *generic* function. This means that it looks at the class of its first argument and arranges that the function actually used is the *method* that matches the class. In Portfolio Probe `deport` has two methods: `deport.randportBurSt` which is for objects created by `random.portfolio` and `deport.portfolBurSt` which is for objects created by `trade.optimizer`. If `deport` is given an object without a class, then an error occurs (because there is not a default method). However, the methods can be used directly.

3.5 Combining Random Portfolio Objects

You may want to combine some random portfolio objects. Suppose you have objects named `rp1`, `rp2` and `rp3` resulting from calls to `random.portfolio`. You would like these to be in one object as they all have the same constraints (or perhaps they have slightly different constraints but you want them all in the same analysis). The `c` function will put them all together:

```
> rp.all <- c(rp1, rp2, rp3)
```

But not all is well:

```
> deport(rp.all)
Error in deport(rp.all) : no applicable method for "deport"
> summary(rp.all)
      Length Class  Mode
[1,] 45      -none- numeric
[2,] 45      -none- numeric
[3,] 45      -none- numeric
[4,] 45      -none- numeric
...
```

Even though `rp.all` is basically correct, it doesn't have the class that the other objects have. Without the class, generic functions like `summary`, `deport` and `valuation` don't work as expected.

```
> class(rp.all) <- class(rp1)
> deport(rp.all)
[1] "rp.all.csv"
```

Once the class is put on the object, we can operate as usual.

Almost. If you want to use `randport.eval`, then you need the `call` attribute as well. In that case, you could give the big object all of the attributes of one of the original objects:

```
> attributes(rp.all) <- attributes(rp1)
```

3.6 Unsatisfiable and Difficult Constraints

Not all sets of constraints can be achieved. Obviously there are no portfolios that satisfy a variance that is smaller than zero (or even smaller than the minimum variance given the other constraints). If you set `random.portfolio` such a task, it is bound to fail.

There is a trade-off between returning quickly when asked the impossible and being successful when asked the merely difficult. There is, of course, a default value for this trade-off, but you can adjust it for specific circumstances. There are a number of control arguments that say how hard to work.

In terms of impossible constraints, the most important is `init.fail`. This says how many separate attempts to make before quitting when there have been no portfolios successfully found.

For each attempt there is a trio of arguments that control how hard to work within the attempt. `iterations.max` gives the maximum number of iterations before stopping. `fail.iter` is the maximum number of consecutive iterations allowed that fail to make progress. `miniter` gives the minimum number of iterations allowed even if `fail.iter` says to quit. (Of course if a portfolio is found that satisfies all the constraints, then the attempt is declared successful and stops no matter what the value of `miniter`.)

The remaining argument of this ilk is `gen.fail`. Let's start with the problem that this argument solves. Suppose you have set a difficult problem for `random.portfolio` and you want 1000 portfolios. Suppose further that the first attempt was successful (so clearly the problem is not impossible) but the next one million attempts fail. As far as you are concerned you are waiting forever. `gen.fail` times the number of portfolios requested is the maximum number of failed attempts allowed.

In our example you requested 1000 and the default value of `gen.fail` is 4, so it would stop after 4000 failures and return the one random portfolio it successfully found (and warn you that it didn't do so well with your request).

Note that it is seldom obvious whether a specific set of constraints is easy to satisfy, difficult to satisfy or impossible.

3.7 Adding a Utility Constraint

The `random.portfolio` function does not allow a constraint on the utility, but `random.portfolio.utility` does. If computation time is of concern, then it can be better to just use `random.portfolio` with constraints on the variance and expected returns. However, this may not be quite what you want.

There is not much difference between the functions in terms of how they are used. The key difference is the `objective.limit` argument. The objective is the negative of the utility. So if you want the utility to be at least 0.6, then you want the argument:

```
objective.limit = -0.6
```

The meaning of `gen.fail` is the same, but the other control arguments are those used with optimization.

There are two forms of failure:

- the objective does not achieve `objective.limit`
- the objective is okay, but there are other broken constraints

The `objfail.max` argument controls how many of the first type are allowed. If `objfail.max=1` (the default) and the first try does not achieve the objective limit, then an error is triggered.

The calls look like:

```
> rp1 <- random.portfolio(100, the.prices, ...)
> ru1 <- random.portfolio.utility(100, -0.6, the.prices, ...)
```

3.8 Going Farther

Tasks that you might want to undertake:

- Chapter 4 discusses how to specify the constraints.
- To review common mistakes, see Section 10.1 on page 121.
- Chapter 5 shows the commands used for examples of random portfolio tasks.

Chapter 4

Constraints

This chapter covers the constraints that can be imposed for generating random portfolios and optimization. `random.portfolio` and `trade.optimizer` use the exact same set of constraint arguments.

4.1 Summary of All Constraints

Table 4.1 lists the possible constraints along with the arguments used to achieve them. In addition to these explicit constraints, there is the implicit constraint of trading integer numbers of asset units.

Round Lots

Trading is only done in integer numbers of asset units except when there is an existing position that is non-integral. Thus if the prices are given for lots as opposed to shares, then round lotting is automatic.

4.2 Monetary Value of the Portfolio

This section discusses the arguments:

- `gross.value`
- `net.value`
- `long.value`
- `short.value`
- `turnover`
- `long.only`
- `allowance`

Table 4.1: Summary of constraints.

Arguments	Constraint	Section
gross.value net.value long.value short.value allowance	monetary value of the portfolio	4.2
turnover	turnover (buys plus sells)	4.2
long.only	no short values if TRUE	4.2
max.weight	maximum weight in the portfolio per asset	4.3
universe.trade	restrict assets to be traded	4.3
lower.trade upper.trade	lower and upper bounds on the the number of asset units to trade	4.3
ntrade	number of assets to trade	4.4
port.size	number of assets in the portfolio	4.4
threshold	threshold constraints on trade and portfolio	4.5
forced.trade	trades that must be done	4.6
positions tol.positions	otherwise available constraints expressed in monetary terms	4.7
lin.constraints lin.bounds lin.trade lin.abs lin.style lin.direction	linear (and/or count) constraints on the portfolio and/or the trade	4.8 4.9
alpha.constraint	bound on expected return of the portfolio	4.10
var.constraint	bound on variance of the portfolio	14.3
bench.constraint	bound on squared tracking error	4.12
dist.center dist.style dist.bounds dist.trade dist.utility dist.coef	distance from one or more portfolios	4.13
sum.weight	maximum of the sum of a specified number of the largest weights	4.14
limit.cost	allowable range of costs	4.15
close.number	number of positions to close	4.16

While there is no single monetary argument that needs to be given, it is mandatory that the monetary value of the portfolio be constrained somehow.

All of the monetary arguments are in the currency units that `prices` uses.

In all cases it is sufficient to only give `turnover`. The argument:

```
turnover = 12000
```

says that the buys plus sells of the trade can not exceed 12,000 currency units.

While this makes most sense when there is an existing portfolio, that is not necessary. The turnover can, of course, be constrained even when other monetary constraints are given.

The turnover can be expressed as an interval:

```
turnover = c(11000, 12000)
```

When only one number is given, the implied lower bound is zero.

How the other monetary arguments are used largely depends on whether or not the portfolio is long-only.

Long-only Portfolios

If you want long-only portfolios, then you need to set the `long.only` argument to `TRUE` (the default is `FALSE`).

You can state the amount of money in the resulting portfolio by giving the `gross.value` argument. Ultimately this needs to be a range of allowable values. You can give the range explicitly with a vector of two numbers:

```
gross.value = c(999900, 1e6)
```

Alternatively you can give a single number:

```
gross.value = 1e6
```

When a single number is given, this is taken to be the upper bound—the lower bound is computed via the `allowance` argument. The default allowance is 0.9999, that is, one basis point away. So (by default) these two specifications of `gross.value` are equivalent.

In general there is no problem with a constraint this tight—the key thing is how wide the range is relative to the prices of the assets. There will be a warning if the interval is seen to be too narrow.

In the case of long-only portfolios, `net.value` and `long.value` are synonyms for `gross.value`, so you can give any one of these three.

Long-short Portfolios

The arguments `gross.value`, `net.value`, `long.value` and `short.value` control the value of the portfolio.

To be clear: The long value is the amount of money in positive positions. The short value is the amount of money in negative positions—this is meant to be a positive number, but the absolute value of negative numbers is taken for

the short value. The gross value is the sum of the long and short values. The net value is the long value minus the short value.

There are two minimal sets of these arguments:

- `gross.value` and `net.value`
- `long.value` and `short.value`

Here are some examples:

```
gross.value = 1e6, net.value = c(200, 3000) # OK
long.value = 6e5, short.value = 5e5 # OK
gross.value = 1e6, net.value = c(0, 2e5), long.value=6e5 # OK
gross.value = 1e6, long.value = 6e5 # not OK, neither pair
```

These four arguments are logically each of length two—giving an allowable range. If they only have length one, then the second value is computed by multiplying the given value by the `allowance` argument. The default value for `allowance` is 0.9999, that is, one basis point away—you may want to alter this depending on how important the tightness of these constraints is to you, and on the size of the portfolio relative to the prices of the assets.

The allowance computation is unlikely to be what is desired for `net.value`, so it is recommended that `net.value` always have length two. Section 7.4 on page 106 gives more details about constraining the value of the portfolio.

There are two cases that are of particular interest: dollar neutral portfolios and portfolios in the genre of 120/20.

Dollar Neutral Portfolios

A dollar neutral portfolio implies that the net value is zero. This is a case where the rule that `net.value` should only be given as an interval *might* be relaxed. The argument:

```
net.value = 0
```

translates into an interval that is symmetric around zero and the radius of the interval is the gross value times one minus `allowance`.

Even so it is probably better to explicitly set your interval for the net value.

120/20 and 130/30 Portfolios

The simplest way to get these is to give constraints just like the stated aim:

```
long.value = 1.2 * NAV, short.value = 0.2 * NAV
```

or possibly a range can be given:

```
long.value = c(1.2, 1.25) * NAV,
short.value = c(0.2, 0.25) * NAV,
net.value = c(.999, 1.0) * NAV
```


4.3 Limits on Assets

This section discusses the arguments:

- `max.weight`
- `enforce.max.weight`
- `universe.trade`
- `lower.trade`
- `upper.trade`

`max.weight`

One typical use of the `max.weight` argument is:

```
max.weight = 0.05
```

This limits the maximum weight of each asset to 5% of the gross value of the portfolio. The weight is the absolute value of the monetary value of the position divided by the gross value. (So even in long-short portfolios the absolute weights always sum to 1.)

The other typical use is to give `max.weight` a named vector that states the maximum weight for each asset. Assets that are not named are not restricted.

Example

If you want to allow a few assets to have larger weights, then you can create a vector of all the pertinent assets. Suppose you want almost all assets to have weight no more than 5% and a few to have weight no more than 7.5%. Then you could create the appropriate vector by:

```
> maxw.spec <- rep(0.05, length=length(prices))
> names(maxw.spec) <- names(prices)
> maxw.spec[spec.assets] <- 0.075
```

The `maxw.spec` vector would then be given as the `max.weight` argument.

A related argument is the control argument `enforce.max.weight`. When this is `TRUE` (the default) and if there are positions in the existing portfolio that will break the maximum weight constraint, then forced trades are created to make them conform to the maximum weight.

caution

There are circumstances in which the `max.weight` argument does not guarantee that the maximum weight in the final portfolio is obeyed if the weight is too large in the existing portfolio. A warning may not be given (since `max.weight` merely limits the extent of trading of the assets).

One case is if the control argument `enforce.max.weight` is `FALSE`. When this argument is `TRUE` (the default), then forced trades are automatically built

to make the positions conform to their maximum weights. This is done for the maximum of the range of the gross value. If the range for the gross value is large and the resulting portfolio has a gross value substantially smaller than the maximum allowed gross, then some maximum weights could be broken.

It is also possible that a maximum weight is broken by enough that the trade can not be forced to be large enough. In this case, the trade is forced to be as large as possible if `enforce.max.weight` is `TRUE`.

If the maximum weight is the same for all assets, then you can ensure that it is obeyed by using `sum.weight` (well, either it's obeyed or you see a warning about it). To just constrain the largest weight, the `sum.weight` argument (see Section 4.14) would look like:

```
sum.weight = c("1"=.1)
```

This restricts the largest weight to 10%.

universe.trade

The `universe.trade` argument should be a vector of character strings containing some of the asset names. Those will be the only assets allowed to trade.

If this argument is `NULL` (the default), then there is no constraint on the assets to trade (except possibly for benchmarks).

lower.trade and upper.trade

The `lower.trade` and `upper.trade` arguments are given in asset units (shares, lots or whatever). As their names suggest, they limit the number of units that can be sold or bought. Like `max.weight` they can be given a single unnamed number or a vector with names corresponding to assets. Assets not named are not restricted.

Values in `lower.trade` can not be positive, and values in `upper.trade` can not be negative. This would be forcing the asset to trade. You can force trades, but not with these arguments—see `forced.trade` (page 54) or possibly `positions` (page 54).

One use of these arguments is to create a selection of assets to buy and a selection of assets to sell as seen on page 103 for long-short portfolios and on page 96 for long-only portfolios. More traditional uses include imposing liquidity constraints and constraining the amount of particular assets in the final portfolio.

Liquidity Constraints

Suppose that you want to constrain trades to a certain fraction of average daily volume. The commands you would use to do this would be similar to:

```
> liquidity <- 0.25 * ave.daily.volume
```

Then you could use this like:

```
upper.trade = liquidity, lower.trade = -liquidity
```

The process of constraining the liquidity merely involves computing the maximum amount of each asset that should be traded; setting the `upper.trade` argument to this amount; and setting the `lower.trade` argument to the negative of that amount.

In this example it is assumed that `ave.daily.volume` is a vector of numbers that has names which are the names of the assets—similar to the `prices` vector. These two vectors need not have the same assets. If there are assets in the daily volume vector (and hence in `liquidity`) that are not in `prices`, then these will be ignored. If `prices` has assets that are not in `liquidity`, then the missing assets will not be limited by the `upper.trade` and `lower.trade` arguments.

caution

Make sure that the units for the volume are the same as those for the prices. An easy mistake would be to try to limit trading to 10% of daily volume, but instead limit it to 10 times daily volume because the volume is in shares while the prices are in lots.

Another sort of liquidity constraint is to limit the amount of assets in the portfolio to n days of average volume. This is more complicated since we want the bounds on the portfolio, but here we have constraints on the trade. This is a case where using `positions` is an easier and safer approach. See page 55 for an example of getting this constraint using `positions`.

4.4 Number of Assets

This section discusses the arguments:

- `ntrade`
- `port.size`

Often these two constraints should be thought of as convenient shortcuts to threshold constraints (page 52).

Number of Assets to Trade

The `ntrade` argument controls the number of assets that may be traded. Generally it is given as a single number, meaning the maximum number to trade. The argument:

```
ntrade = 25
```

says that no more than 25 assets will be traded.

If you do not want a limit on the number of assets traded, then set `ntrade` to the size of the universe (or larger). The default is to have no limit.

If you desire a minimum number of assets traded, then give `ntrade` a length two vector. For example:

```
ntrade=c(4, 25)
```

states that the number of assets traded needs to be between 4 and 25, inclusive. A minimum number to trade is most likely to be useful when threshold constraints (Section 4.5) are used as well. Otherwise trading just one unit of an asset counts.

The minimum number to trade is more useful for generating random portfolios than in optimizing.

Number of Assets in the Portfolio

The `port.size` argument constrains the number of assets in the constructed portfolio. Like `ntrade`, this can be given as either a single number (meaning the upper bound), or as a length 2 vector giving the range.

The argument:

```
port.size = 50
```

means no more than 50 assets may be in the portfolio, while:

```
port.size = c(50, 50)
```

means that exactly 50 assets need to be in the portfolio.

4.5 Threshold Constraints

This section discusses the argument:

- `threshold`

The `threshold` argument controls two types of threshold constraints—trade thresholds and portfolio thresholds. Threshold constraints may also be specified with `positions` (page 54). Note that in any one problem you can only declare threshold constraints with one of these arguments (though both arguments can be used in one problem).

Trade Thresholds

Trade threshold constraints force trades to be at least a certain amount of an asset if it is traded at all.

The argument:

```
threshold=c(ABC=4, DEF=23)
```

requests that at least 4 units (lots perhaps) of asset ABC and at least 23 units of DEF be bought or sold if they are traded at all. This is not the same as forcing a trade of a certain size—that is discussed in Section 4.6.

When the `threshold` argument is a vector (or a one-column matrix), then the constraint is taken to be symmetric. Another way of stating the constraint given above would be:

```
threshold=rbind(ABC=c(-4,4), DEF=c(-23,23))
```

where what we are giving to `threshold` looks like:

```
> rbind(ABC=c(-4,4), DEF=c(-23,23))
      [,1] [,2]
ABC    -4    4
DEF   -23   23
```

Give a two-column matrix when you want different constraints for buying and selling for at least one of the assets.

```
threshold=rbind(ABC=c(-5,4), DEF=c(0,23))
```

The command above states that if ABC is sold, then at least 5 units should be sold. If ABC is bought, then at least 4 units should be bought. If DEF is bought, then at least 23 units should be bought. There is no threshold constraint if DEF is sold.

Portfolio Thresholds

A portfolio threshold constraint states the minimum number of units of an asset that should be held in the portfolio. For example, you may prefer to have no lots of ABC if less than 7 lots are held. Portfolio thresholds are specified similarly to trade thresholds except they are in the third and fourth columns of the matrix instead of the first and second.

Here is an example:

```
> thresh.m1 <- cbind(0, 0, rbind(ABC=c(-7, 7), DEF=c(-5, 8)))
> thresh.m1
      [,1] [,2] [,3] [,4]
ABC     0    0   -7    7
DEF     0    0   -5    8
```

If `thresh.m1` is given as `threshold`, then there should be at least 7 units of ABC—either long or short—if it is in the portfolio at all. Also DEF should be at least 5 short or at least 8 long or not in the portfolio. In this case there are no trading threshold constraints since the first two columns are all zero.

Summary of Threshold Inputs

- A vector or a one-column matrix: symmetric trading constraints.
- A two-column matrix: symmetric or asymmetric trading constraints.
- A three-column matrix: first two columns are trading constraints, third column is portfolio constraint for long-only portfolios. It is an error to give a three-column matrix with long-short portfolios.
- A four-column matrix: first two columns are trading constraints, third and fourth columns are portfolio constraints.

4.6 Forced Trades

This section discusses the argument:

- `forced.trade`

The `forced.trade` argument is a named vector giving one or more trades that must be performed. The value gives the minimal amount to trade and the names give the assets to be traded. For example:

```
forced.trade = c(ABC=-8, DEF=15)
```

says to sell at least 8 units (lots) of ABC and buy at least 15 units of DEF. If you wanted to buy exactly 15 units of DEF, then you would say:

```
forced.trade = c(ABC=-8, DEF=15), upper.trade=c(DEF=15)
```

Trades may also be automatically forced if the existing portfolio breaks maximum weight constraints—see page 49). The `positions` argument can also force trades.

4.7 Positions

This section discusses the arguments:

- `positions`
- `tol.positions`

The `positions` argument does not embody any constraints that can't be achieved with other arguments. It exists because the constraints can sometimes be more conveniently expressed via `positions`.

The constraints that `positions` can do are:

- `max.weight`
- `universe.trade`
- `lower.trade`
- `upper.trade`
- `threshold`
- `forced.trade`

The key difference between `positions` and these other constraints is that `positions` is expressed in monetary value. In examples we will assume the currency unit is dollars, but it is really the currency unit that the `prices` vector is in.

The `positions` argument takes a matrix with rows representing assets, and 2, 4 or 8 columns. Not all assets in the problem need to be present in `positions` but (by default) there will be a warning if not.

Portfolio constraints

The first two columns of `positions` contain portfolio constraints. Column 1 is the minimum amount of money that is allowed in the portfolio for each of the assets. Column 2 is the maximum amount of money.

So this includes the `max.weight` constraint. If both `positions` and `max.weight` are given, then whichever is the stronger for each asset is the actual constraint.

n days of daily volume in portfolio

On page 50 there is an example of imposing liquidity constraints on the trade. Here we want to impose liquidity constraints on the portfolio.

Suppose that we want to restrict positions in the portfolio to be no more than 8 days of average volume. Since `positions` is expecting monetary value rather than the number of asset units, we first need to transform to money:

```
> liquid.value <- ave.daily.volume[names(prices)] * prices
> any(is.na(liquid.value)) # trouble if this is TRUE
```

There are a couple things to note here. We are assuming that the volume and the prices refer to the same asset units—we don't want one to be in shares and the other in lots. We'll see on page 56 why the occurrence of missing values would be upsetting.

Now we are ready to build our matrix to give to the `positions` argument. If we have a long-only portfolio, then the matrix can be:

```
> posmat1 <- cbind(0, 8 * liquid.value)
```

If we have a long-short portfolio, then it would be:

```
> posmat2 <- 8 * cbind(-liquid.value, liquid.value)
```

Let's do one more supposition. Suppose that we have a long-only portfolio and we want no position to be larger than \$300,000 as well as having the liquidity constraint. You might be tempted to impose that with a command like:

```
> posmat1[, 2] <- min(3e5, posmat1[, 2]) # WRONG
```

Note that there is no warning from this—S has no way of knowing that you are doing something you don't really want to do. What you do want to do is:

```
> posmat1[, 2] <- pmin(3e5, posmat1[, 2]) # right
```

The `min` function returns just one number which is the minimum of all the numbers it sees. The `pmin` function does a minimum for each element, which is what we want in this instance.

Trade constraints

Trade constraints are imposed with the third and fourth columns of `positions`. If you want to impose trade constraints without imposing portfolio constraints, then you can just make the first two columns infinite. For example, if we want selling of each asset to be limited to \$5000 and buying of each asset to be limited to \$8000, then we could create a matrix like:

```
> posmat3 <- cbind(rep(-Inf, length(prices)), Inf,
+                 -5000, 8000)
> head(posmat3)
      [,1] [,2] [,3] [,4]
[1,] -Inf  Inf -5000 8000
[2,] -Inf  Inf -5000 8000
[3,] -Inf  Inf -5000 8000
[4,] -Inf  Inf -5000 8000
[5,] -Inf  Inf -5000 8000
[6,] -Inf  Inf -5000 8000
> dimnames(posmat3) <- list(names(prices), NULL)
```

These constraints are essentially identical to `lower.trade` and `upper.trade` except they are expressed in money rather than asset units, and it is possible to impose forced trades with the `positions` constraints.

Forced constraints

A trade is forced either if the range for the portfolio does not include the current position, or if the trading range does not include zero. This substitutes for using `forced.trade`.

Forced trades from `positions` cause the output of `trade.optimizer` to have a component named `positions.forced` which will show the forced trades created.

Universe constraints

If there is a missing value (`NA`) in any of the first four columns of `positions`, then the corresponding asset is not allowed to trade. (It is an error to have a missing value in the fifth through eighth columns.) These constraints could be achieved by using the `universe.trade` argument.

Be careful not to let missing values stray in by accident. The `summary` of optimized portfolios tells how many assets `positions` forces not to trade.

Threshold constraints

The 5th through the 8th columns of `positions` are for threshold constraints. All four of these columns need to be given if any are given. Note also that if these columns are given, then the `threshold` argument can not be given.

You can put null constraints in columns that you don't want to have influence. The most extreme possibility in this regard is if you want to use `positions` to only impose a portfolio threshold on a long-only portfolio. Suppose you want only positions that have at least \$20,000, then you could build the `positions` matrix like:


```

> posmat4 <- cbind(rep(-Inf, length(prices)), Inf, -Inf, Inf,
+ 0, 0, 0, 2e4)
> head(posmat4)
      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8]
[1,] -Inf  Inf -Inf  Inf   0   0   0 20000
[2,] -Inf  Inf -Inf  Inf   0   0   0 20000
[3,] -Inf  Inf -Inf  Inf   0   0   0 20000
[4,] -Inf  Inf -Inf  Inf   0   0   0 20000
[5,] -Inf  Inf -Inf  Inf   0   0   0 20000
[6,] -Inf  Inf -Inf  Inf   0   0   0 20000
> dimnames(posmat4) <- list(names(prices), NULL)

```

Tolerance

Imagine a situation where asset ABC is restricted to be no more than \$50,000 in the portfolio. At time 1 the optimizer likes ABC and buys enough that it is close to that upper limit. From time 1 to time 2 ABC dutifully does well, so now it is above the \$50,000 limit. Assuming the limit remains the same for time 2, then the optimizer will force a sale of some ABC. So we are performing a trade (probably of a trivial amount) for a rather arbitrary reason.

If there is a trade threshold constraint on selling ABC, then it is even worse—we are effectively forcing a sale of at least the threshold amount.

The `tol.positions` argument solves this problem. If `tol.positions` is zero (the default), then the constraints are taken as stated. If `tol.positions` is positive, then assets that violate portfolio constraints in positions by a value less than `tol.positions` are not forced to trade.

In our example if the value of the position of ABC at time 2 is \$50,203, then at least \$203 worth of ABC will be forced to be sold if `tol.positions` is less than \$203. But there will not be a forced trade if `tol.positions` is greater than \$203.

Note that this tolerance only applies to assets that are already in the existing portfolio.

Summary of positions inputs

- 2-column matrix: lower and upper constraints on the value of portfolio positions.
- 4-column matrix: 1-2) constraints on portfolio positions; 3-4) constraints on the value of sells and buys in the trade.
- 8-column matrix: 1-2) constraints on the portfolio positions; 3-4) constraints on the trade; 5-6) threshold constraints on the values of trades; 7-8) threshold constraints on the values in the portfolio. (Note: portfolio-trade-trade-portfolio.)

4.8 Linear Constraints

This section describes the arguments:

- `lin.constraints`
- `lin.bounds`
- `lin.trade`
- `lin.abs`
- `lin.style`
- `lin.direction`

These arguments create linear constraints on the portfolio and/or the trade. These arguments can also impose count constraints—a topic that is discussed in the next section.

Portfolio Probe makes a distinction between numeric constraints and constraints based on categories. An example of the first is a bound on twice the value of asset `ABC` minus 1.5 times the value of asset `DEF` minus the value of asset `GHI`. An example of the second is bounds on the value from each country in the asset universe. (Most other optimizers will expand such constraints into a matrix of zeros and ones—you could do this also, but the method used here is more compact as well as easier for the user.)

Linear constraints are used when there is aggregation across assets. If a constraint only involves one asset at a time, then there are other constraints that are more expedient—for instance, `max.weight` or `threshold`.

Building Constraints

To impose linear constraints you must specify at least two arguments—you can accept default values of other arguments. These two key arguments are:

- `lin.constraints`
- `lin.bounds`

The `lin.constraints` argument gives the numbers or categories for each asset for each constraint. `lin.bounds` provides the lower and upper bound for each (sub)constraint. Because the functions are picky about these arguments, it is best to create them with the `build.constraints` function.

The `build.constraints` function takes a vector, a matrix, a factor or a data frame containing the information for the `lin.constraints` argument. The vector or matrix can be character, numeric or logical. A data frame is required if you have a mixture of types.

In this example, we are giving a character matrix with two columns—one for countries and one for sectors:

```
> cons.obj <- build.constraints(cbind(country=countryvec,
+   sector=sectvec))
> names(cons.obj)
[1] "lin.constraints" "bounds"
```

The result of `build.constraints` is a list with two components. The first component, `lin.constraints`, is suitable to give as the `lin.constraints` argument, the `bounds` component is the template for an object to give as `lin.bounds`. The `lin.constraints` object must have column names, so `build.constraints` will add column names if they are not already there. (The column names are used to keep track of which bounds go with which constraint.)

```
> cons.obj$lin.constraints[1:3,]
      country sector
ABC "Spain"  "media"
DEF "France" "retail"
GHI "Italy"  "energy"
> cons.obj$bounds
      lower upper
country : France -Inf  Inf
country : Italy  -Inf  Inf
country : Spain  -Inf  Inf
sector  : energy -Inf  Inf
sector  : media  -Inf  Inf
sector  : retail -Inf  Inf
sector  : telecom -Inf  Inf
```

The typical use of the `bounds` component is to create a separate object out of it, and then change any of the infinite values desired.

Bounds and `lin.style`

What the bounds look like depends on the style that you pick for them. Bounds on linear constraints may be in terms of either weight or monetary value. The `lin.style` argument controls this— a value of `"weight"` versus a value of `"value"`. `lin.style` also accepts a third value (`"count"`) which is discussed on page 67.

If we want bounds in terms of weights (the default), then we could create the bounds matrix something like:

```
> bounds.csw <- cons.obj$bounds
> bounds.csw[1:2, ] <- c(.1, .2, .4, .5)
> bounds.csw[3, 2] <- .2
> bounds.csw[5, ] <- c(.3, .6)
> bounds.csw
      lower upper
country : France  0.1  0.4
country : Italy   0.2  0.5
country : Spain  -Inf  0.2
sector  : energy -Inf  Inf
sector  : media  0.3  0.6
sector  : retail -Inf  Inf
sector  : telecom -Inf  Inf
```

We could then use the arguments:

```
lin.constraints = cons.obj$lin.constraints,
lin.bounds = bounds.csw, gross.value = 1e6
```

If on the other hand, we want bounds in terms of monetary value, then we could do:

```
> bounds.csv <- cons.obj$bounds
> bounds.csv[1:2, ] <- c(1e5, 2e5, 4e5, 5e5)
> bounds.csv[3, 2] <- 2e5
> bounds.csv[5,] <- c(3e5, 6e5)
> bounds.csv
              lower upper
country : France 1e+05 4e+05
country : Italy  2e+05 5e+05
country : Spain  -Inf  2e+05
sector  : energy -Inf   Inf
sector  : media  3e+05 6e+05
sector  : retail -Inf   Inf
sector  : telecom -Inf   Inf
```

Now the arguments would be:

```
lin.constraints = cons.obj$lin.constraints,
lin.bounds = bounds.csv, gross.value = 1e6,
lin.style = "value"
```

In this case we need to give the `lin.style` argument because we are using a non-default value for it.

It is also possible to mix styles:

```
> bounds.csmix <- bounds.csw
> bounds.csmix[1:3, ] <- bounds.csv[1:3, ]
> bounds.csmix
              lower upper
country : France 1e+05 4e+05
country : Italy  2e+05 5e+05
country : Spain  -Inf  2e+05
sector  : energy -Inf   Inf
sector  : media   0.3  0.6
sector  : retail -Inf   Inf
sector  : telecom -Inf   Inf
```

So the arguments become:

```
lin.constraints = cons.obj$lin.constraints,
lin.bounds = bounds.csv, gross.value = 1e6,
lin.style = c("value", "weight")
```

There is no need for any particular pattern of finite bounds. In this example, the only sector we are bounding is media.

Once a bounds matrix has been set up, it can be used when building new constraint objects. Suppose that we want to change from sectors to industries, we can build new constraints like:

```
> cons.obj2 <- build.constraints(cbind(country=countryvec,
+   industry=industvec), bound=bounds.csw)
```

S code note

In the `cbind` command, we are assuming that `countryvec` and `industvec` have the same assets in the same order. A safer approach would be:

```
> ci.inam <- intersect(names(countryvec), names(industvec))
> cbind(country=countryvec[ci.inam],
+   industry=industvec[ci.inam])
```

The bounds component of `cons.obj2` will have the same bounds for the countries as `bounds.csw` and will have infinite bounds for the industries.

A bounds object that is actually used may contain extraneous bounds—for example bounds for sectors when only countries are being constrained. However, it is an error not to give bounds for all of the constraints represented in `lin.constraints`.

Only categorical constraints have been discussed so far. We now look at numerical constraints.

Numerical Constraints: Risk Factors

Constraining risk factors is perhaps the most common numerical linear constraint. In a fixed income portfolio duration would be an example.

Whether constraints are numerical or categorical, the first step is to use `build.constraints`:

```
> cons.obj3 <- build.constraints(cbind(Risk1=beta1))
> head(cons.obj3$lin.constraints)
      Risk1
stockA 1.3119461
stockB 0.9886379
stockC 1.1688637
stockD 0.8160228
stockE 1.0312180
stockF 1.2067453
> cons.obj3$bounds
      lower upper
Risk1  -Inf   Inf
> cons.obj3$bounds[] <- c(0.9, 0.95)
> cons.obj3$bounds
      lower upper
Risk1   0.9   0.95
```

S code note

Two things:

First, the use of `cbind` in the first line of the example above is merely an easy way to create a one-column matrix that has a name for the column.

Second, you might be tempted to do:

```
> cons.obj3$bounds <- c(0.9, 0.95)
```

instead of

```
> cons.obj3$bounds[] <- c(0.9, 0.95)
```

The difference is that the first of these will result in a plain vector while the second retains the matrix structure of the original object. We want that matrix structure.

We would now use the arguments:

```
lin.constraints = cons.obj3$lin.constraints,
lin.bounds = cons.obj3$bounds
```

Constraining more than one risk factor is straightforward.

```
> cons.obj4 <- build.constraints(cbind(Risk1=beta1,
+ Risk2=beta2))
> head(cons.obj4$lin.constraints)
      Risk1 Risk2
stockA 1.3119461 0.8972907
stockB 0.9886379 0.9784834
stockC 1.1688637 0.7135961
stockD 0.8160228 1.1195767
stockE 1.0312180 1.0847337
stockF 1.2067453 1.0977745
```

Now just create the bounds that you want and use the arguments as before.

Numerical Constraints: Market Capitalization

If you want to have a portfolio that has a similar average market capitalization as an index, then you can impose this as a linear numerical constraint. We assume here that the index is weighted by market capitalization and we have the weights for the index.

The first step is to use `build.constraints` as before:

```
> cons.mcap <- build.constraints(cbind(Mcap=index.wts))
```

The possibly surprising part is what the bound looks like. The value for this constraint (assuming a weight style as opposed to value) that the index has is the sum the squared weights:

```
> mcap.target <- sum(index.wts^2)
```

The bounds you use might, for example, be 95% and 105% of the target.

Mixing Numerical and Categorical Constraints

There is one tricky part when you have both numerical and categorical linear constraints. Up until now we have used a matrix (or a vector) as the first argument to `build.constraints`. A matrix has to have all of its elements of the same type, but now we want a mix of types: numeric for the numerical constraints, and character, logical or factor for the categorical constraints. Now we need to use a data frame, not a matrix, to represent our constraint data.

```
> cons.objmix <- build.constraints(data.frame(Risk1=beta1,
+ Risk2=beta2, Country=country))
> head(cons.objmix$lin.constraints)
      Risk1 Risk2 Country
stockA 1.3119461 0.8972907 Australia
stockB 0.9886379 0.9784834 Singapore
stockC 1.1688637 0.7135961 Japan
stockD 0.8160228 1.1195767 Australia
stockE 1.0312180 1.0847337 Japan
stockF 1.2067453 1.0977745 Australia
> cons.objmix$bounds
              lower upper
Risk1          -Inf   Inf
Risk2          -Inf   Inf
Country : Australia -Inf   Inf
Country : Japan     -Inf   Inf
Country : Singapore -Inf   Inf
```

Yet again the next step is to specify the bounds and then use the arguments.

Portfolio Constraints versus Trade Constraints

The default behavior is to constrain the portfolio. If you want the trade to be constrained rather than the portfolio, then set the `lin.trade` argument to `TRUE`. You can have a mixture of constraints on the trade and on the portfolio. The vector given as `lin.trade` is replicated to have length equal to the number of columns in the `lin.constraint` object.

If `lin.constraints` has three columns, then the argument:

```
lin.trade = c(TRUE, FALSE, TRUE)
```

means that the constraint in the first column is on the trade, the second column constrains the portfolio and the third column constrains the trade.

Net Constraints versus Gross Constraints

By default, constraints use the sum of weights (or value or counts). If you want the constraints on the gross (that is, the absolute value) rather than the net, then set the `lin.abs` argument to be `TRUE`.

For example, in a long-short portfolio you may want to limit the net exposure to a sector, but you may also want to limit the gross amount in the sector.

The `lin.abs` argument is replicated to have length equal to the number of columns in the `lin.constraint` object.

As with similar arguments, `lin.abs` can be a vector:

```
lin.abs = c(TRUE, TRUE, FALSE)
```

Let's look at the four possible combinations of the values of `lin.trade` and `lin.abs`. These are:

- portfolio, net (not absolute) (the default)
- portfolio, gross (absolute)
- trade, net (not absolute)
- trade, gross (absolute)

For long-only portfolios the first two are the same since all positions are positive.

An extreme example is to enforce all four types on the same constraint. You want to name the constraints properly so that you can keep track of them:

```
> con.4c <- build.constraints(cbind(c.tn=countryvec,
+   c.ta=countryvec, c.pn=countryvec, c.pa=countryvec))
```

The result of the above command can then be used like:

```
lin.constraints = con.4c$lin.constraints,
lin.trade = c(TRUE,TRUE,FALSE,FALSE),
lin.abs = c(FALSE,TRUE))
```

The `lin.abs` argument will automatically be replicated to be: FALSE, TRUE, FALSE, TRUE. Of course we could have given all four values.

Long-side Constraints and Short-side Constraints

This section applies to long-only portfolios for trade constraints but not portfolio constraints. For long-short portfolios it applies to both trade constraints and portfolio constraints.

A long-side constraint only involves assets that have long positions—the short positions are effectively treated as zero for the purposes of the constraint. Likewise, a short-side constraint only looks at assets with short positions.

The `lin.direction` argument controls whether you have a long-side, short-side or both-side constraint.

There are three possible values:

- `lin.direction = 0` means both-sides (the default)
- `lin.direction = 1` means long-side
- `lin.direction = -1` means short-side

Of course, `lin.direction` gets replicated to be as long as the number of columns in the `lin.constraints` object, and can be given as a vector:

```
lin.direction = c(1, -1, 0, 0)
```

When a short-side constraint is in effect (`lin.direction=-1`), then it is the absolute value of the short values that is used—your constraints should be non-negative.

Looking at the Effect of the Constraints

The `constraints.realized` function identifies the status of linear constraints for a portfolio. This function produces the `con.realized` component of objects returned by `trade.optimizer`. This component is not a part of the printing of the object, but is a part of the summary. (Using `summary` can give you information on some of the other constraints as well.) Here is an example:

```
> opti2$con.realized
$linear
      lower upper realized nearest violation
c.tn : France -3e+05 -2e+05 -296632 -3368 NA
c.tn : Italy 1e+05 2e+05 180767 19233 NA
c.tn : Spain -1e+05 -3e+04 -75494 -24506 NA
c.ta : France 6e+05 7e+05 600624 -624 NA
c.ta : Italy 2e+05 3e+05 206917 -6917 NA
c.ta : Spain 4e+05 5e+05 399552 NA -448
c.pn : France -1e+05 0e+00 -24099 24099 NA
c.pn : Italy 0e+00 1e+05 92422 7578 NA
c.pn : Spain -1e+05 0e+00 -76340 -23660 NA
c.pa : France 3e+05 4e+05 394621 5379 NA
c.pa : Italy 3e+05 4e+05 304686 -4686 NA
c.pa : Spain 3e+05 4e+05 400398 NA 398
```

This contains a matrix where each row has information on a sub-constraint. The final column indicates the size of violations—if all of the elements in this column are `NA`, then there are no violations. A negative violation means that the realized value is below the lower bound, and a positive violation means that the realized value is above the upper bound. The next to last column gives proximity of the realized value to the nearest bound (if there is not a violation). The same convention in terms of sign is used—a negative number means the realized is closer to the lower bound. The last two columns have exactly one missing value for each row.

In this example the sizes of the violations are relatively trivial. However, they would not be trivial in terms of their affect on the objective if you were optimizing—any violation means that the optimizer has been concentrating on getting a feasible solution (a solution that merely satisfies the constraints) and may not be near the best feasible solution. At times there is a non-zero penalty that is of trivial size relative to the utility, you need not worry in such cases.

Evaluating Un-imposed Constraints

It is possible to look at the value of constraints that were not imposed on a portfolio. Give `constraints.realized` the portfolio object (the result of a call to `trade.optimizer`) and the constraints that are of interest:

```
> constraints.realized(opt.ir2, cbind(country=countryvec))
              lower upper realized nearest violation
country : France -Inf  Inf  0.020106    -Inf      NA
country : Italy  -Inf  Inf  0.156728    -Inf      NA
country : Spain  -Inf  Inf -0.157020    -Inf      NA
```

You may also use the additional linear constraint arguments like `lin.abs` and `lin.trade` to specify the nature of the constraints you want:

```
> constraints.realized(opt.ir2, cbind(country=countryvec),
+   lin.abs=TRUE, lin.style="value")
              lower upper realized nearest violation
country : France -Inf  Inf  433494    -Inf      NA
country : Italy  -Inf  Inf  296702    -Inf      NA
country : Spain  -Inf  Inf  669666    -Inf      NA
```

Inspecting Linear Constraints

The `summary` for portfolios includes the result of `constraints.realized` on the portfolio as well as listing the effective values of the arguments that control the type of constraints. If you are doing optimization, you can get this with:

```
opt.obj <- trade.optimizer( ... )
summary(opt.obj)
```

If you are generating random portfolios, you can do:

```
rp.obj <- random.portfolio( ... )
randport.eval(rp.obj, subset=1, FUN=summary)
```

Setting the `subset` argument to 1 means that we are only going to evaluate the first random portfolio.

An example for random portfolios is:

```
> tail(randport.eval(rp.obj6, FUN=summary,
+   subset=1)[[1]])
$valuation.trade.fraction.of.gross
      gross      net      long      short
2.00683544 0.01830419 1.01256982 0.99426563
$constraints.realized
$constraints.realized$linear
              lower upper realized nearest violation
Country : France   -300    0   -28.09   28.09      NA
```

```

Country : Germany      0  100   27.57  -27.57      NA
Country : UK          100 2000  175.39  -75.39      NA
$lin.trade
[1] FALSE
$lin.abs
[1] FALSE
$lin.style
[1] "value"
$lin.direction
[1] 0

```

An example for optimization is:

```

> tail(summary(op.obj6))
$valuation.trade.fraction.of.gross
      gross      net      long      short
0.05179854 0.02020499 0.03600176 0.01579677
$constraints.realized
$constraints.realized$linear
      lower  upper realized  nearest violation
Cntry : France -300.0   0.00 -297.6500 -2.350e+00      NA
Cntry : Germany  0.0 100.00  11.4600 -1.146e+01      NA
Cntry : UK      100.0 2000.00 782.1500 -6.822e+02      NA
Beta           0.9   0.95   0.9001 -1.060e-04      NA
$lin.trade
[1] FALSE
$lin.abs
[1] FALSE
$lin.style
  Cntry  Beta
"value" "weight"
$lin.direction
[1] 0

```

Notice that if a linear argument is the same for all constraints, then just one value is printed. On the other hand if an argument has more than one value (`lin.style` in the example above), then all values are printed— including names of the constraints.

4.9 Count Constraints

This section continues the description of arguments:

- `lin.constraints`
- `lin.bounds`
- `lin.trade`
- `lin.abs`
- `lin.style`

- `lin.direction`

but with the restriction that `lin.style="count"`.

When `lin.style` is "count", then count constraints—user-defined integer constraints— are produced. This can only be used with a categorical column of `lin.constraints`—a count constraint on a numerical column will trigger an error.

Let's consider some examples.

Suppose that we have a long-short portfolio and we want an equal number of long positions and short positions from the UK, and we want one more French long position than French short position. We have no preference regarding Germany. Then we would create a bounds object like:

```
> lbc1
                lower upper
Country : France    0.5   1.5
Country : Germany -Inf   Inf
Country : UK       -0.5   0.5
```

Our call would then include the arguments:

```
lin.bounds = lbc1, lin.style = "count"
```

We can accept the default values of `lin.abs` and `lin.trade` because we want a net constraint on the portfolio. We can infer from the bounds object that the `lin.constraints` argument contains just one constraint that is called `Country`.

Notice that in this context (`lin.abs = FALSE`, `lin.direction = 0`) long positions count as one but short positions count as negative one. Also notice that we are giving bounds that are between integers—you don't want to give bounds for count constraints that are integer.

Now suppose that we want exactly one German position in the portfolio and we want one or two French positions. Our bounds object should now look like:

```
> lbc2
                lower upper
Country : France    0.5   2.5
Country : Germany    0.5   1.5
Country : UK       -Inf   Inf
```

We would have arguments:

```
lin.bounds = lbc2, lin.style = "count", lin.abs = TRUE
```

We are adding `lin.abs` to the call because we want an absolute constraint.

If we want exactly two long UK positions and we want no short French positions, then we need to do a bit more work. The extra work is because we now have two constraints (a long-side constraint and a short-side constraint) so we need two columns in `lin.constraints`, not the one column that we have been using. We start by building the appropriate object:

```
> lcc3 <- build.constraints(cbind(Country.longc=countryvec,
+   Country.shortc=countryvec))
```

Now we get the bounds matrix that we need:

```
> lbc3 <- lcc3$bounds
> # lbc3
> lbc3[3,] <- c(1.5, 2.5)
> lbc3[4, 2] <- .5
> lbc3
```

	lower	upper
Country.longc : France	-Inf	Inf
Country.longc : Germany	-Inf	Inf
Country.longc : UK	1.5	2.5
Country.shortc : France	-Inf	0.5
Country.shortc : Germany	-Inf	Inf
Country.shortc : UK	-Inf	Inf

Now we'll use arguments:

```
lin.constraints = lcc3$lin.constraints, lin.bounds = lbc3,
lin.style = "count", lin.direction = c(1, -1)
```

Logical values are allowed to define constraints. Logicals are probably most likely to be used in count constraints. Here we have a somewhat contrived example (as if the others aren't) where we want to lower the beta of the portfolio. We lower beta if we sell high beta stocks or buy low beta stocks. We'll create a constraint that does more of this than its opposite in terms of asset count. We build the constraint object and the bounds:

```
> lcc4 <- build.constraints(cbind(HighBeta = beta > 1)
> lbc4 <- lcc4$bounds
> lbc4[1,1] <- 3.5
> lbc4
```

	lower	upper
HighBeta : FALSE	3.5	Inf
HighBeta : TRUE	-Inf	Inf

Our arguments are:

```
lin.constraints = lcc4$lin.constraints, lin.bounds = lbc4,
lin.style = "count", lin.trade = TRUE
```

4.10 Alpha (Expected Return) Constraints

This section describes the argument:

- `alpha.constraint`

The `alpha.constraint` argument bounds the expected return of the optimized portfolio. In its simplest use, it is merely a single number that gives the minimum for the expected return. For example:

```
alpha.constraint = 1.1
```

An upper bound as well as a lower bound can be specified by giving a two-column matrix. For instance:

```
alpha.constraint = cbind(1.1, 2.3)
```

will restrict the expected return to be between 1.1 and 2.3.

Note that if a benchmark is given, then this will constrain the expected return relative to the benchmark. This is just the usual expected return of the portfolio minus the expected return of the benchmark.

Advanced use of `alpha.constraint` (chiefly when more than one vector of expected returns is given) is discussed on page 147.

4.11 Variance Constraints

This section discusses the argument:

- `var.constraint`

The `var.constraint` argument constrains the value of the variance. In its most common usage, it is merely a number giving the upper bound for the variance. Using the argument:

```
var.constraint = 1.4
```

imposes an upper bound of 1.4 on the variance.

The next most common usage is a two column matrix that gives a range of values that the variance is allowed to be in. For example:

```
var.constraint = cbind(1.2, 1.4)
```

says that the variance is to be at least 1.2 but no more than 1.4. The value given looks like:

```
> cbind(1.2, 1.4)
      [,1] [,2]
[1,] 1.2  1.4
```

More advanced use of this argument—predominantly when multiple variances are given—is discussed on page 145.

4.12 Tracking Error (Benchmark) Constraints

This section discusses the argument:

- `bench.constraint`

The `bench.constraint` argument is used to constrain (squared) tracking errors. There has to be information on not only the numerical bound that you want, but the benchmark as well. The benchmark needs to be an asset in the variance, but it doesn't need to be in the prices unless benchmark trading is allowed.

Ultimately `bench.constraint` creates a variance constraint—it is just a more convenient way of specifying some commonly desired constraints.

Single Upper Bound

If you want just an upper bound, then give a named numeric vector of length one. For example:

```
bench.constraint = c(spx = .04^2)
```

specifies a 4% tracking error from `spx`.

Another way of doing the same thing is:

```
> bc <- .04 ^ 2
> names(bc) <- "spx"
> bc
  spx
0.0016
```

and then use the argument:

```
bench.constraint = bc
```

Scaling

Don't forget that this constraint is on squared tracking error, not tracking error.

While tracking errors are almost always an annual number, your data probably aren't annual. For example if your variance is on a daily scale, then you need to transform the constraint to the daily scale:

```
bench.constraint = c(spx = .04^2 / 252)
```

Lower and Upper Bounds

Lower bounds as well as upper bounds can be specified by giving a two-column matrix. For example:

```
bench.constraint = rbind(spx=c(.02, .04)^2)
```

would restrict `spx` tracking error to between 2% and 4%.

This is most likely of interest when generating random portfolios.

Multiple Benchmarks

If you want upper bounds on more than one benchmark, merely give a numeric vector with the names of the benchmarks:

```
bench.constraint = c(spx = .04^2, R1000 = .05^2)
```

specifies no more than a 4% tracking error from `spx` and also no more than a 5% tracking error from `R1000`.

To get both upper and lower constraints for multiple benchmarks, you give a two-column matrix with multiple rows. For example:

```
bench.constraint = mult.benchcon.lowup
```

where

```
> mult.benchcon.lowup
      [,1] [,2]
spx    0.0004 0.0016
R1000  0.0009 0.0025
R2000  0.0016 0.0064
```

Advanced Use

When there are multiple variance matrices, then specifying benchmark constraints is more complicated. This is discussed on page 145.

4.13 Distance

This section discusses the arguments:

- `dist.center`
- `dist.style`
- `dist.bounds`
- `dist.trade`
- `dist.utility`
- `dist.prices`

These arguments impose distance constraints.

Let's start with a simple example. Suppose we have a target portfolio that has the following weights:

```
> targwt <- c(Asset1=.4, Asset2=.3, Asset3=.2, Asset4=.1)
```


Suppose we want the resulting portfolio to have weight distance no more than 0.5 (so the buys plus sells of the trade to get from the result to the target would be no more than 50% of the gross value of the portfolio). To impose that constraint we would use arguments:

```
dist.center=targetwt, dist.style="weight", dist.bounds=.5
```

The example above computes distance as the difference in weights. The distance can alternatively be computed as the difference in monetary values. These two are equivalent if the gross value is constrained to a narrow range (although the weight formulation takes slightly more time to compute). However, if the gross value is allowed a wide range, then there is a difference: the value distance will favor a gross value close to the gross value of the target portfolio while the weight distance will be indifferent to the gross value.

For a value distance, you can give either the values of the assets in the target portfolio or the shares in the target portfolio. Using value might look like:

```
dist.center=target.values, dist.style="value",
dist.bounds=25000
```

while using shares might look like:

```
dist.center=target.shares, dist.style="shares",
dist.bounds=25000
```

The `dist.trade` argument takes a logical vector and allows you to constrain the distance of the trade rather than the portfolio.

The `dist.utility` argument also takes a logical vector—TRUE values of this imply that you want to use it as the objective in optimization. For more on this, see Section 6.4.

Alternative prices

By default, distance is defined by the `prices` vector. You don't have to define it that way. You can include a value for `dist.prices` that gives different prices to be used in the computations of distances. The `dist.style` argument takes a value with “custom” before the style when you want to use the `dist.prices` values. For example, you might have:

```
dist.center=target.shares, dist.style="customshares",
dist.prices=some.price.vector, dist.bounds=35000
```

Multiple distances

It is possible to have more than one distance constraint. In this case the `dist.center` argument needs to be a list. For example, you might have arguments that look like:

```
dist.center=list(target1.wt, target2.share, target3.wt),
dist.style=c("weight", "shares", "weight"),
dist.bounds=c(0.2, 2e5, 0.3)
```

Here we have three distances declared—the first and third in terms of weight and the second in terms of shares.

If you want lower bounds as well as upper bounds on the distances, then you need to give `dist.bounds` a two-column matrix with each row corresponding to a distance. For example:

```
dist.bounds=cbind(c(0.1, 1e5, 0.2), c(0.2, 2e5, 0.3))
```

Each of the arguments `dist.style`, `dist.trade` and `dist.utility` can either have length equal to the number of distances, or length one if the argument is the same for all arguments. For example, in the example above `dist.trade` and `dist.utility` both have their length one default value.

The `dist.prices` argument can be a list of price vectors to match the `dist.center` argument.

4.14 Sums of Largest Weights

This section discusses the argument:

- `sum.weight`

Consider the constraint that the largest n weights should sum to no more than some limit—this is handled by the `sum.weight` argument.

For example if you want the 5 largest weights to sum to no more than 40%, and the 10 largest weights to sum to no more than 70%, the argument would look like:

```
sum.weight = c("5"=.4, "10"=.7)
```

The values in the vector given to `sum.weight` are the limits, the names of the vector are the numbers of weights in the sums.

S code note

The names within the `c` function must be inside quotes as they are not legal S object names.

Another way of getting the same thing would be:

```
> sumw.arg <- c(.4, .7)
> names(sumw.arg) <- c(5, 10)
```

Then `sumw.arg` would be given as the `sum.weight` argument.

The argument:

```
sum.weight = c("1" = .05)
```

creates a constraint that says that the largest weight of any one asset is 5%. This is conceptually exactly the same as:

```
max.weight = .05
```

However, there are some operational differences. The key difference is that using `sum.weight` ensures that either the constraint is satisfied or there is a penalty imposed for breaking the constraint. But it will be somewhat slower when `sum.weight` is used. See page 49 for the possibility of `max.weight` failing to enforce the constraint.

4.15 Cost Constraints

This section discusses the argument:

- `limit.cost`

This is unlikely to be of use with optimization.

`limit.cost` puts lower and upper constraints on the cost. This can be useful for mimicking actual trading with random portfolios. If this is given, then it must be a length two vector giving the allowable range of costs:

```
limit.cost = c(2.3, 2.35)
```

or possibly something like:

```
limit.cost = opt.obj$results["cost"] * c(.99, 1.01)
```

4.16 Number of Positions to Close

This section discusses the argument:

- `close.number`

The minimum and maximum number of positions to close in the existing portfolio can be specified with the `close.number` argument. If a single number is given, then exactly that number are to be closed.

So:

```
close.number = 3
```

is really shorthand for:

```
close.number = c(3, 3)
```

Note that this is different than other arguments where a single number does not imply a lower bound.

To specify that at least 2 positions should be closed, but no more than 3, then say:

```
close.number = c(2, 3)
```

This sort of constraint is unlikely to be useful in optimization, but can be used to generate random portfolios that match what an actual trade has done.

4.17 Quadratic Constraints

Quadratic constraints on the portfolio are not officially supported, but can be imposed. Quadratic constraints on the trade are not possible.

The acceptance of more than one variance means that `trade.optimizer` and `random.portfolio` can be tricked into allowing quadratic constraints on the portfolio.

Here are the steps for creating quadratic constraints:

- Add the constraint matrices to the variance.
- Specify the constraint bounds.
- Make sure the constraints are not benchmark-relative.
- If doing optimization, tell the optimizer not to include the constraints in the utility.

Add Constraints to the Variance

Suppose you have a single variance matrix (as per usual) called `varian` and you have two quadratic constraint matrices called `qmat1` and `qmat2`. You need to create a three dimensional array with three slices. This is an easy operation in the S language, but care needs to be taken that the assets match.

```
> assetnam <- dimnames(varian)[[1]]
> varian3 <- array(c(varian, qmat1[assetnam, assetnam],
+   qmat2[assetnam, assetnam]), c(dim(varian), 3),
+   c(dimnames(varian), list(NULL)))
```

S code note

The `c` function features prominently in this command.

When `c` is used on matrices, as in its first occurrence in the command, the dimensions are stripped and the result is a plain vector of all of the elements of the three matrices.

The `c` function can be used with lists as well as with atomic vectors. For the `dimnames` of the resulting three-dimensional array, we need to have a length three list where the final component is ignored (and hence can be `NULL`). In this case we need all of the arguments to `c` to be lists, so we need to give it `list(NULL)`. Just giving an argument of `NULL` would not have had the effect that we wanted.

Impose Constraint Bounds

The next step is to impose the constraint bounds. This will be done with the `var.constraint` argument as discussed in Section 14.3. The constraints in this case will be on the second and third variances, that is, on indices 1 and 2. Also note that the bounds are in terms of the weights—there is not a choice to put the bounds in terms of monetary value.

If you are generating random portfolios (without a utility constraint) and you are not using a benchmark, then you can proceed to the actual computation. The intervening steps only apply if you are using a benchmark or a utility.

Dummy Run

The next two steps are aided by performing a dummy run of optimization—something like:

```
> qc.temp <- trade.optimizer(prices, varian3,
+   iterations=0, ...)
```

It is best to include the variance constraints in this run—without them the step regarding benchmarks can go wrong. When `iterations.max` is set to zero, then a minimal amount of optimization is performed. You will get a warning that this is not the same as no optimization. We don't care about optimization in this instance, we only care about the setup for optimization.

Check for Benchmark

If you are not using a benchmark, you don't need to be concerned with this step.

If you are using a benchmark, then the default behavior makes the variances relative to the benchmark. This is seen in the `vtable` component of the output of our dummy run.

```
> qc.temp$vtable
      [,1] [,2] [,3]
variance    0    1    2
benchmark  500  500  500
utility.only  1    0    0
attr("benchmarks")
[1] "spx" "spx" "spx"
```

In our example we want the last two columns to not have a benchmark. Hence we would do an operation like:

```
> qcvtab <- qc.temp$vtable
> attr(qcvtab, "benchmarks") <- c("spx", "", "")
```

If you are generating random portfolios and you didn't include the variance constraints in the dummy run, then you will also want to change the `utility.only` row to be zero in the columns that pertain to your constraints. The argument:

```
vtable = qcvtab
```

should be given in the actual computation.

You can learn more about the `vtable` argument on page 144.

Constraints out of Utility

This step does not pertain to random portfolios (unless a utility constraint is used).

This step tells the optimizer that the final two “variances” are not to be used in the utility. We do this by modifying the utility table from the dummy run. (Utility tables are discussed on page 149.)

```
> qc.utiltab <- qc.temp$utable
> qc.utiltab
      [,1] [,2] [,3]
alpha.spot      0   0   0
variance.spot   0   1   2
destination     0   1   2
opt.objective   1   1   1
risk.aversion   1   1   1
wt.in.destination 1   1   1
> qc.utiltab['destination', ] <- c(0, -1, -1)
> qc.utiltab
      [,1] [,2] [,3]
alpha.spot      0   0   0
variance.spot   0   1   2
destination     0  -1  -1
opt.objective   1   1   1
risk.aversion   1   1   1
wt.in.destination 1   1   1
```

Actual Computation

Now we are ready for the actual computation. If we are generating random portfolios, we’ll have a command like:

```
> qc.rp <- random.portfolio(100, prices, varian3,
+   vtable = qcvtab, var.con = ...)
```

Alternatively, if we are optimizing, we’d do something like:

```
> qc.opt <- trade.optimizer(prices, varian3,
+   vtable = qcvtab, utable = qc.utiltab, var.con = ...)
```

4.18 Constraint Penalties and Soft Constraints

This section applies only to optimization—it can not apply to generating random portfolios.

Virtually all of the constraints are enforced via penalties. That is, instead of the optimization algorithm minimizing the negative utility, it minimizes the negative utility plus penalties for whatever constraints are violated. Solutions that obey all of the constraints experience no penalties.

The algorithm actively attempts to satisfy some of the constraints—in particular the constraints on the monetary value of portfolios are closely managed,

as are threshold constraints, and the constraint on the maximum number of assets traded is always obeyed.

Tools are available in `trade.optimizer` to allow you to adjust the penalties so that most of the constraints (those not actively managed) can become soft constraints. This is done by making the penalties for those constraints small enough so that there will be a trade-off between the penalty for the constraint and the utility.

Consider an example:

```
> soft.op1$violated
[1] "variance"      "linear"        "gross value"
attr(,"linear violations")
[1] "Country" "Sector"
```

There are violations of two linear constraints, the variance constraint, and the gross value.

The default value of the `penalty.constraint` argument is 1000. You can change individual elements of the penalty by giving the `penalty.constraint` argument a named vector where the names are the items for which you want to change the penalty. You need to give the names exactly—they can not be abbreviated. The easiest approach is to change the `penalty.constraint` component of the output from the original optimization.

```
> softpen2 <- soft.op1$penalty.constraint
> softpen2[1:2] <- .1
> soft.op2 <- update(soft.op1, penalty.constraint=softpen2)
```


Chapter 5

Details of Random Portfolio Examples

This chapter outlines the code that underlies the examples given in Chapter 2. The files that contain the exact commands are available on the website. There are at least two uses for this:

- You can reproduce the analyses in Chapter 2.
- If you have a task similar to one of the examples, you should be able to do what you need to do by mimicking what is done here.

The analyses are somewhat simplified—they ignore such things as dividends, stock splits and other corporate actions. The analyses also depend on some custom functions that can be found in file `pprobe_functions01.R`.

A set of adjusted closing prices on a non-random selection of large US stocks was retrieved from the internet. The set of commands for creating the data is given in file `pprobe_R_data01.txt`.

The first step is to read in the set of identifiers:

```
> us.selection <- scan("us_selection.txt", what="")
```

The function used next depends on the TTR package [Ulrich, 2009]. If you don't already have it installed on your machine, then you need to do:

```
> install.packages("TTR")
```

Some of these instructions may need to be changed if you are using S-PLUS rather than R. The command above is one of those.

The next step is to get the price data:

```
> us.price <- pp.TTR.multsymbol(us.selection, 20060101,  
+ 20090101)
```

The price data are then used to calculate the MACD matrix:

```
> us.macd <- pp.TTR.multmacd(us.price)
```

This function merely calls the (univariate) `MACD` function from `TTR` multiple times and saves the *signal* part of the output.

The objects `us.price` and `us.macd` are of class `xts`. In order to avoid assuming that `xts` is available, we convert these to matrices:

```
> us.pricemat <- as.matrix(us.price)
> us.macdmat <- as.matrix(us.macd)
```

(The only substantial change is that row names are created for the matrices.)

R Commands to Do It All

This describes one way of using the auxiliary files to reproduce the analyses. We assume that there is a directory called `topdir`. This has a subdirectory called `test`. If the variance caching is done, it will have another subdirectory called `Varshr`.

The auxiliary files from the website are placed in `topdir`. R is started in `test`. The scripts include an object called `PP_LOC` that you can change if you would like a different topology for the directories.

In the R session start off by making sure you are in the right place. The result of:

```
> getwd()
```

should be `".../topdir/test"`.

Then you want to install the regular functions and the graphics functions:

```
> source("../pprobe_functions01.R")
> source("../pprobe_graphFun01.R")
```

Now you need to build the basic data:

```
> source("../pprobe_R_data01.txt")
```

This retrieves data from the internet and builds the basic data objects. This will fail if the first asset ID is not available, but should work if others are missing. (This currently produces data for 69 stocks rather than the original 70—`BNI` has been dropped.)

If you are going to use the cached variance matrices, then execute the file that contains those commands:

```
> source("../pprobe_R_varCache.txt")
```

The static method of performance measurement is done with:

```
> source("../pprobe_R_perfMeas_static.txt")
```

To test constraint bounds, do:

```
> source("../pprobe_R_bounds.txt")
```

Risk is explored with:

```
> source("../pprobe_R_risk.txt")
```

The shadowing method of performance measurement is done with:

```
> source("../pprobe_R_perfMeas_shadow.txt")
```

Backtests are evaluated with:

```
> source("../pprobe_R_varCache.txt")
```

5.1 Performance Measurement: Static

The file `pprobe_R_perfMeas_static.txt` gives the full set of commands for these examples.

Step number one is to create the initial portfolio. This was done by optimizing a portfolio:

```
> init.20.w10 <- trade.optimizer(us.pricemat[251,],
+   expected.return=us.macdmat[251,], port.size=20,
+   max.weight=.1, gross.value=1e6, long.only=TRUE
+   utility="maximum return", do.warn=c(novar=FALSE))
```

The 200% turnover path of optimal portfolios was created with:

```
> ser.20.w10.t200 <- pp.serial.opt(init.20.w10,
+   pricemat=us.pricemat[-1:-251,], alphamat=us.macdmat,
+   frac.turnover=2/252, max.weight=.1, port.size=20,
+   do.warn=c(novar=FALSE, utility.switch=FALSE),
+   keep.port=TRUE)
```

An object with 400% turnover was also created.

Now, we generate random portfolios that obey the constraints at the start of the time period:

```
> rp251.20.w10 <- random.portfolio(1000, us.pricemat[251,],
+   port.size=20, max.weight=.1, gross.value=1e6,
+   long.only=TRUE)
```

We then calculate the random portfolio returns over the period:

```
> rp251.20.w10.ret07 <- pp.simpret(valuation(rp251.20.w10,
+   us.pricemat[c(251, nrow(us.pricemat)-252)],
+   collapse=TRUE))
```

We also need to calculate the return of the optimal fund throughout 2007:

```
> serret200.07 <- ser.20.w10.t200$value[251] /
+   ser.20.w10.t200$value[1] - 1
```

For this one time period and amount of turnover, the only thing left to do is the comparison. The percentile can be computed as:

```
> round(100 * mean(rp251.20.w20.ret07 < serret200.07))
```

A visual representation can be done like:

```
> plot(density(100 * rp251.20.w20.ret07))
> abline(v = 100 * serret200.07, col="blue")
```

If you want to save plots, then you can open an appropriate graphics device before you perform the graphics commands and close the graphics device afterwards. For example:

```
> pdf(file="myplots.pdf")
> plot(density(100 * rp251.20.w20.ret07))
> abline(v = 100 * serret200.07, col="blue")
> dev.off()
```

5.2 Testing the Effect of Constraint Bounds

The complete set of commands is given in `pprobe.R.bounds.txt`.

Static Utility

First, we create a price vector at a specific time:

```
> price.20061229 <- us.pricemat[251, ]
```

We create a return matrix from the price matrix:

```
> us.retmat <- diff(log(us.pricemat))
```

We also need a specific variance matrix:

```
> varshr.Q1 <- var.shrink.eqcor(us.retmat[1:250,])
```

The next step is to generate random portfolios with all the constraints except the volatility constraint:

```
> rp30w10 <- random.portfolio(number.rand = 10000,
+   prices = price.20061229, gross.value = 1e+06,
+   long.only = TRUE, port.size = 30, max.weight = 0.1)
```

The utility (with risk aversion 2) of these portfolios can now be computed for the year 2007:

```
> rp30w10.mvu2.2007 <- pp.meanvarutil(valuation(rp30w10,
+   us.price[252:502,], collapse=TRUE), 2)
```

The same two steps (generate portfolios, calculate utility) are now done with the volatility constraint:

```

> rp30w10v08 <- random.portfolio(number.rand = 10000,
+   prices = price.20061229, variance=varshr.Q1,
+   gross.value = 1e+06, long.only = TRUE, port.size = 30,
+   max.weight = 0.1, var.constraint=.08^2/252)
> rp30w10v08.mvu2.2007 <- pp.meanvarutil(valuation(rp30w10v08,
+   us.price[252:502,], collapse=TRUE), 2)

```

Caching Variances

Many operations depend on a series of variance matrices. Variance matrices are sometimes reasonably large so the memory used by a few hundred of them can become a problem. One solution is to create a directory containing individual files of each variance that can each be loaded as needed. The solution shown here works for R, it would (probably) need to be adjusted some to work in S-PLUS. The code that creates the files of variances may be found in file `pprobe_R_varCache.txt`.

Here is code to create the files:

```

> for(i in 250:700) {
+   JJd <- rownames(us.retmat)[i]
+   JJn <- paste("varshr", JJd, sep=".")
+   assign(JJn, var.shrink.eqcor(us.retmat[seq(to=i,
+     length=250),]))
+   save(list=JJn, file=paste("Varshr/", JJn, ".rda",
+     sep=""))
+   rm(list=JJn)
+   cat("done with", JJd, "\n")
+ }

```

The following fragment from a function shows how the variance files can be used:

```

for(i in seq(along=dates)) {
  this.date <- dates[i]
  vnam <- paste(varprefix, this.date, sep="")
  attach(paste(vardir, "/", vnam, ".rda", sep=""))
  this.rand <- random.portfolio(
    variance=get(vnam), ...)
  detach()
  ....
}

```

In each iteration the file containing the desired variance matrix is attached, the variance is used, and then the file is detached. The function has a couple of arguments (`vardir` and `varprefix`) that allow some flexibility.

Dynamic Utility

The first dynamic utility object does not involve the variances. The following call sets the window size to 60 trading days and the risk aversion to 2, among other things:

```

> rp30w10.mvu2.dat <- pp.date.meanvarutil(
+   rownames(us.pricemat)[251:695], 60, 2,
+   us.pricemat, us.macdmat, max.weight=.1,
+   port.size=30, gross.value=1e6, long.only=TRUE,
+   vardir=paste(PPLOC, "Vardir", sep="/"))

```

Even though the variances aren't really used, we still need to provide a valid location for the directory holding the variances because the function still pulls the variances in.

The next command finds the minimum variance that can be achieved with the given constraints. We don't really care if we have the exact minimum or not, so in the interests of speed the number of iterations that the optimizer is allowed is reduced to 3:

```

> rp30w10.minvar.dat <- pp.date.minvar(
+   rownames(us.pricemat)[251:695], us.pricemat,
+   max.weight=.1, port.size=30, gross.value=1e6,
+   long.only=TRUE, iteration=3,
+   vardir=paste(PPLOC, "Varshr", sep="/"))

```

The command to create the dynamic utility with the volatility constraint is:

```

> rp30w10vv20.mvu2.dat <- pp.date.var.meanvarutil(
+   rownames(us.pricemat)[251:695],
+   1.44 * rp30w10.minvar.dat, 60, 2, us.pricemat,
+   us.macdmat, max.weight=.1, port.size=30,
+   gross.value=1e6, long.only=TRUE, throw.error=FALSE,
+   vardir=paste(PPLOC, "Varshr", sep="/"))

```

We want a constraint that is 120% of the minimum volatility, but the constraint is actually on the variance—the square of 120% is 1.44.

5.3 Test a Risk Model

The full set of commands for this task is in file `pprobe_R_risk.txt`.

A preliminary step is to create a vector stating where in the price data the quarters end:

```

> quarter.ends <- c(251, 312, 375, 438, 502, 563,
+   627, 691, 755)
> names(quarter.ends) <- paste(rep(2006:2008,
+   c(1,4,4)), 'Q', c(4,1:3), sep="")

```

We create the template for a number of objects that we'll fill in:

```

> empty8 <- vector("list", 8)
> names(empty8) <- tail(names(quarter.ends), 8)

```

We make a return matrix using the price matrix:

```

> us.retmat <- diff(log(us.pricemat))

```

Ex Ante Volatility Comparisons

We load the package that contains the variance estimators (installation of this package is described on page 110):

```
> library(BurStFin)
```

Now we initialize lists to hold the variances and then estimate the variances:

```
> varshr.Q <- empty8
> varfac.Q <- empty8
> for(i in 1:8) {
+   varshr.Q[[i]] <- var.shrink.eqcor(
+     us.retmat[seq(to=quarter.ends[i]-1,
+       length=250),])
+   varfac.Q[[i]] <- factor.model.stat(
+     us.retmat[seq(to=quarter.ends[i]-1,
+       length=250),])
+ }
```

Next, we create a set of random portfolios for each quarter:

```
> rp30s15.Q <- empty8
> for(i in 1:8) {
+   rp30s15.Q[[i]] <- random.portfolio(1000,
+     us.pricemat[quarter.ends[i], ], gross.value=1e6,
+     long.only=TRUE, port.size=30, max.weight=.05,
+     sum.weight=c("4"=.15))
+ }
```

What we want from the random portfolios is the predictions of their volatility. So we create an object to hold the predictions from each type of model:

```
> rp30s15.facexante.Q <- empty8
> rp30s15.shrexante.Q <- empty8
> for(j in 1:8) {
+   rp30s15.facexante.Q[[j]] <- unlist(randport.eval(
+     rp30s15.Q[[j]], additional.args=list(prices=
+     us.pricemat[quarter.ends[j], ],
+     variance=varfac.Q[[j]]), keep="var.values"),
+     use.names=FALSE)
+   rp30s15.shrexante.Q[[j]] <- unlist(randport.eval(
+     rp30s15.Q[[j]], additional.args=list(prices=
+     us.pricemat[quarter.ends[j], ],
+     variance=varshr.Q[[j]]), keep="var.values"),
+     use.names=FALSE)
+ }
```

We generally think in terms of volatility, but it is actually (daily) variance that is saved into the objects.

advanced

The commands above ensure that a subtle but important problem is avoided. The `randport.eval` function looks at the image of how the random portfolio object was created (the `call` attribute) and uses that to create the actual data objects that are used. The `prices` vector for the random portfolios was:

```
us.pricemat[quarter.ends[i], ]
```

In the current loop there is no `i`, so either an arbitrary `i` will be used or there will be an error if no `i` is found. We could have had our loop variable be `i`, which would have been fine in this case, but there is no guarantee that the loop variables correspond to the same thing.

There is a check that will warn you if it detects that the current data are different from the original. If the data are largely static, then the check can fail to properly detect the problem.

How the current loop has avoided the problem is to explicitly state the `prices` vector and variance matrix that we want to use by putting them in the `additional.args` argument.

We can compare the volatility estimates in a plot like:

```
> plot(100 * sqrt(252 * rp30s15.facexante.Q[[4]]),
+      100 * sqrt(252 * rp30s15.shrexante.Q[[4]]))
```

Ex Ante versus Realized Volatility

We compute the realized volatility of the random portfolios over each of the quarters:

```
> rp30s15.realized.Q <- empty8
> for(i in 1:8) {
+   rp30s15.realized.Q[[i]] <- pp.realvol(
+     valuation(rp30s15.Q[[i]],
+     us.pricemat[ quarter.ends[i]:quarter.ends[i+1], ],
+     collapse=TRUE))
+ }
```

To make a plot comparing a prediction and the realized, do:

```
> plot(100 * sqrt(252 * rp30s15.facexante.Q[[4]]),
+      rp30s15.realized.Q[[4]])
> abline(0, 1, col="blue")
```

We now compute correlation between predicted and realized volatility:

```
> corshrreal <- numeric(8)
> names(corshrreal) <- names(rp30s15.realized.Q)
> corfacreal <- corshrreal
> for(i in 1:8) {
```



```

+ corshrreal[i] <- cor(rp30s15.shrexante.Q[[i]],
+   rp30s15.realized.Q[[i]], method='spearman')
+ corfacreal[i] <- cor(rp30s15.facexante.Q[[i]],
+   rp30s15.realized.Q[[i]], method='spearman')
+ }

```

Actually this is computing the correlation between the predicted variance and the realized volatility. Since we are using Spearman (rank) correlations, it doesn't matter whether we use variance or volatility—the ordering is unchanged when we convert from variance to volatility.

Next we bootstrap the correlations and compute 95% confidence intervals:

```

> corboot.Q <- rep(list(numeric(5000)), 8)
> names(corboot.Q) <- names(empty8)
> for(i in 1:5000) {
+   this.bootsamp <- sample(1000, 1000, replace=TRUE)
+   for(j in 1:8) {
+     corboot.Q[[j]][i] <- cor(rp30s15.shrexante.Q[[j]][
+       this.bootsamp], rp30s15.realized.Q[[j]][
+       this.bootsamp], method='spearman')
+   }
+ }
> corshrreal.95ci <- sapply(corboot.Q, quantile,
+   probs=c(.025, .975))

```

The confidence intervals aren't all that important, but they provide a nice example of showing how easy it is to perform statistical bootstrapping.

The dollar neutral and 120/20 portfolios are done analogously. The key difference is in the call to `valuation` to get the realized volatility. The `type` and `cash` arguments become important—the assumption of how much cash there is makes a difference to the returns and hence to the volatility.

5.4 Performance Measurement: Shadowing

The commands for this section are in file `pprobe_R_perfMeas_shadow.txt`.

The first thing is to create the optimized run as shown in Section 5.1.

Then the random runs that imitate the optimal run are done:

```

> randser.20.w10.t200 <- pp.serial.rand(200,
+   init.20.w10, pricemat=us.pricemat[-1:-251,],
+   alphamat=us.macdmat, frac.turnover=2/252,
+   max.weight=.1, port.size=20, do.warn=c(novar=FALSE,
+   utility.switch=FALSE))

```

A plot can now be performed:

```

> the.den <- density(100 * pp.simpret(
+   randser.20.w10.t200$value[c(1,252),]))
> plot(the.den, main="", xlab="2007 return (percent)",
+   xlim=range(the.den$x, 100 * serret200.07))
> abline(v=100 * serret200.07, col="blue", lwd=2)

```

A percentile can also be calculated:

```
> round(mean(pp.simpret(randser.20.w10.t200$value[c(1,252),])
+       < serret200.07) * 100)
```

5.5 Evaluating a Trading Strategy

The code for this section is in file `pprobe_R_tradeStrat.txt`.

The first step is to generate a set of random portfolios (20 of them in this case) to serve as the starting points:

```
> rp.20.w10.2008 <- random.portfolio(20,
+   us.pricemat[502,], gross.value=1e7, long.only=TRUE,
+   max.weight=.1, port.size=20)
```

Doing the optimization and then the random mimicking for each of the starting portfolios is all done in the following function:

```
> evalst.20.w10.t200.2008 <- pp.evalstrat(100,
+   rp.20.w10.2008, us.pricemat[-1:-502,],
+   us.macdmat, frac.turnover=2/252, max.weight=.1,
+   port.size=20, verbose=2,
+   do.warn=c(utility.switch=FALSE, novariance=FALSE))
```

The percentiles for a point in time can be plotted:

```
> plot(ppoints(20) * 100,
+   sort(evalst.20.w10.t200.2008$percentile[61,]),
+   ylim=c(0,100), xlab="Theoretical quantiles",
+   ylab="Percentile")
> abline(0,1, col="blue")
```

You can see what such a plot would like with a purely random strategy with commands like:

```
> plot(ppoints(20) * 100, sort(sample(0:100, 20,
+   replace=TRUE)))
> abline(0, 1, col="blue")
```

5.6 Going Farther

- You can modify the scripts to use your own data.
- You can modify the functions that the scripts use.

Chapter 6

Optimizing Long-Only Portfolios

This chapter discusses commands for optimizing long-only portfolios. It applies to both active and passive portfolios.

The examples in this chapter are just sketches of what you are likely to want to do in practice. Chapter 4 discusses all the constraints that may be imposed. Some of the examples limit the number of assets that are held in the portfolio or traded. These limits are convenient, but threshold constraints (page 52) may be better to use. Chapter 9 covers trading costs—it is usually a good idea to constrain trading in some way, whether with trading costs or a turnover constraint.

6.1 Required Inputs

The `trade.optimizer` function needs a price vector. The `prices` argument needs to be a vector of prices with names that are the asset names. The assets in `prices` control the assets that are used in the problem. Other inputs—such as the variance and linear constraints—may contain additional assets (which will be ignored).

Other than `prices` there is no single argument that is required. However, there are two concepts that are required: the amount of money in the portfolio has to be constrained somehow, and there has to be something that provides a utility.

Monetary Value

The turnover or the amount of money in the portfolio needs to be constrained. Details can be found on page 45

Utility

At least one of the arguments:

- `variance`

- `expected.return`
- `dist.center`

must be given.

Utility-free optimization—minimizing the distance to a target portfolio—can be done using `dist.center` and related arguments. See Section 6.4 to do this.

Assuming distance is out of the picture: If `expected.return` is not given, then the utility is the minimum variance. If a variance matrix is not given, then the utility is the maximum expected return. If both are given, then the default utility is to maximize the information ratio. It can be safest to explicitly state the utility that you want.

Optimizing a long-only portfolio without a variance or a distance is very unlikely to be a good idea.

6.2 Examples for Passive Portfolios

This section contains some examples that build in complexity. Expected returns do not enter into optimization for passive portfolios, and they need not be given.

Minimize the Variance of the Portfolio

This is the simplest optimization to perform:

```
> op.mv1 <- trade.optimizer(prices, varian,
+   utility = "minimum variance",
+   long.only=TRUE, gross.value=1.4e6)
```

The amount put into the portfolio is 1.4 million currency units. The simplicity of this command means that it may not do what you want since it is dependent on default values for most arguments.

Minimize Tracking Error

We can minimize the variance relative to a benchmark—that is, minimize tracking error. To be concrete, let's assume that the benchmark is the S&P 500 with the asset name of `"spx"`.

The benchmark needs to be in the variance matrix. If it isn't and you have the weights of the constituents, refer to page 112. The benchmark need not be in the price vector unless trading the benchmark is allowed (via the `bench.trade` argument).

A simple version of minimizing the tracking error is:

```
> op.te1 <- trade.optimizer(prices, varian,
+   utility = "minimum variance",
+   long.only=TRUE, gross.value=1.4e6, benchmark="spx")
```

This is almost surely too simple.

For example, you may want to limit the number of assets in the portfolio:

```
> op.te2 <- trade.optimizer(prices, varian,
+   utility = "minimum variance",
+   long.only=TRUE, gross.value=1.4e6, benchmark="spx",
+   port.size=55, max.weight=0.05)
```

The `op.te2` object represents a portfolio containing up to 55 assets that (approximately) minimizes the tracking error among all portfolios of size 55 in which assets have a maximum weight of 5%.

If you are rebalancing an existing portfolio rather than creating a new one, then the command might be:

```
> op.te3 <- trade.optimizer(prices, varian,
+   utility = "minimum variance",
+   long.only=TRUE, gross.val=1.4e6, benchmark="spx",
+   max.weight=0.05, existing=cur.port, ntrade=20)
```

The command creating `op.te3` is allowing up to 20 assets to be traded. It does not control the turnover, nor the number of assets in the resulting portfolio. Thus `op.te3` can potentially have up to 75 assets in it if the current portfolio contains 55 assets. Suppose that we want to have a portfolio with 45 to 55 assets in it, and we want to hold the turnover to 100,000 (currency units, e.g., dollars). The following command will provide that:

```
> op.te4 <- trade.optimizer(prices, varian,
+   utility = "minimum variance",
+   long.only=TRUE, gross.value=1.4e6, benchmark="spx",
+   max.weight=0.05, existing=cur.port, turnover=1e5,
+   port.size=c(45,55))
```

This command is a realistic possibility.

The examples above could have been done without specifying the `utility` argument—since there are no expected returns, the utility would default to minimum variance (but you would get a warning stating that utility was being modified).

See page 98 for examples of dealing with cash flow. Page 150 has an example of minimizing against dual benchmarks—for example, when you have a prediction of the benchmark after an upcoming rebalancing.

6.3 Examples for Active Portfolios

Expected returns of the assets are required for optimization of active portfolios. (This is not absolutely true—see page 96 for an example.) Additionally, there needs to be a decision on how to balance expected return with risk.

Maximize the Information Ratio

The easiest optimization that involves both expected returns and variance is to maximize the information ratio. The information ratio of a portfolio is the expected return of the portfolio divided by the standard deviation of the portfolio. (The standard deviation is the square root of the variance.) This is a

natural quantity (in mean-variance land) to maximize—we're seeking the best risk-adjusted return.

A command to maximize the information ratio is:

```
> op.ir1 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6)
```

The `alphas` object is a vector of expected returns of the assets. When both variance and expected returns are given, the default objective is to maximize the information ratio. In this example 1.4 million currency units are put into the portfolio.

The information ratio is traditionally put into annualized terms. If the variance and expected returns in your command are not annualized, then you can annualize the information ratio by multiplying the utility by the square root of the number that you would use to annualize the variance and expected returns. For example, for a daily frequency the annualized information ratio would be:

```
> -op.ir1$utility.value * sqrt(252)
```

The negative sign is because the optimizer always minimizes. When quantities are being maximized (information ratio or utility), then it really minimizes the negative of the quantity.

Note that this is the information ratio you would get if all of your inputs were exact. The actual information ratio that you achieve is likely to be worse.

We now turn to the case of maximizing the information ratio when there is a benchmark involved.

The Information Ratio with a Tracking Error Constraint

It is common to perform a mean-variance optimization relative to the benchmark. The risk aversion is adjusted so that the tracking error is near its desired level. It makes more financial sense to maximize the information ratio in absolute terms (as in the command that created `op.ir1`) while having a constraint that the tracking error is no larger than a given value.

The mechanism in `trade.optimizer` that allows a constraint relative to a benchmark is the `bench.constraint` argument. Here is a command to constrain the tracking error to be no more than 4% while maximizing the information ratio:

```
> op.ir2 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6, bench.constraint = c(spx=0.04^2/252))
```

The value of the `bench.constraint` argument needs to be named (so that it knows what benchmark you want) and the value is in the scale of the variance. Using the `c` function allows a name to be put on the value.

The 0.04 is squared to transform the tracking error into a variance. This is then divided by 252 in order to go from the annual value to daily, which is the frequency of the variance we are assuming in this example. If the variance were

annualized and created from returns in percent, then the `bench.constraint` argument would have been:

```
c(spx=4^2)
```

note

The `benchmark` argument is not given in the previous command. If it had been, then the information ratio that was optimized would be relative to the benchmark.

A more realistic command might be:

```
> op.ir3 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6, turnover=1e5, ntrade=15,
+   bench.constraint = c(spx=0.04^2/252), max.weight=.05,
+   existing=cur.port, port.size=60)
```

This command includes the existing portfolio, has a limit on the turnover, and specifies that the final portfolio should contain no more than 60 assets.

You can easily add constraints to other benchmarks as well—just make the `bench.constraint` argument longer:

```
> op.ir4 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6, turnover=1e5, ntrade=15,
+   bench.constraint = c(spx=0.04^2/252, newspx=0.05^2/252),
+   max.weight=.05, existing=cur.port, port.size=60)
```

In this example we are constraining the tracking error to be no more than 4% against the current benchmark, and no more than 5% against our prediction of the benchmark after it is rebalanced.

Minimizing relative to two or more benchmarks is, however, somewhat more complicated—that is discussed on page 150. Also a little more complicated is to get a benchmark-relative utility but a constraint on the portfolio variance—page 152 shows how to do that.

Maximize Benchmark-relative Information Ratio

If you want the the utility—in this case the information ratio—to be relative to a benchmark, then you merely need to give the `benchmark` argument:

```
> op.ir5 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6, benchmark = "spx")
```

You get substantially different behavior depending on whether you use `benchmark` or `bench.constraint`:

- `bench.constraint` does not affect which utility is used.

- `benchmark` changes the utility.

You can use both—in which case `benchmark` changes the utility as usual, but there is also a constraint on the tracking error.

Mean-Variance Optimization

The default objective is to maximize an information ratio (whenever both variance and expected returns are given). If you want to perform a mean-variance optimization instead, then you need to specify the `utility` argument. The generic mean-variance problem is to maximize the expected return of the portfolio minus the risk aversion times the variance of the portfolio. Page 140 has more on the risk aversion parameter.

A command to do a simple mean-variance optimization is:

```
> op.mv1 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6, utility='mean-var', risk.aversion=.7)
```

This command does not include a benchmark. As noted above, if the `benchmark` argument is given, then the mean and variance will be relative to the benchmark. If you do have a benchmark, an alternative approach is to do the optimization in the absolute sense and add a constraint on the tracking error. The tracking error constraint is placed on the optimization with a command like:

```
> op.mv2 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   gross.value=1.4e6, utility='mean-var', risk.aversion=.7,
+   bench.constraint = c(spx=0.04^2/252))
```

For an explanation of the value given for `bench.constraint`, see page 71.

In mean-variance optimization, the risk aversion parameter is quite important.

Mean-Volatility Optimization

Mean-volatility optimization is just like mean-variance optimization except that the value of the `utility` argument is:

```
utility = "mean-volatility"
```

and the `risk.aversion` argument will (usually) have a substantially different value.

Buy-Hold-Sell List

It is not absolutely necessary to create expected returns for an actively managed portfolio. Instead you can minimize the variance or tracking error using a list of assets that can be sold and a list of assets that can be bought.

Suppose that `buy.vec` and `sell.vec` are vectors giving the assets that are allowed to be bought and sold. The first thing to do is to set up the constraints that make sure nothing on the buy list can be sold, and nothing on the sell list can be bought.


```

> lower.tr <- rep(0, length(buy.vec))
> names(lower.tr) <- buy.vec
> upper.tr <- rep(0, length(sell.vec))
> names(upper.tr) <- sell.vec

```

S code note

These commands create two new vectors containing all zeros and with names that are the buy assets for one of them, and names that are the sell assets for the other. The `rep` function replicates its first argument.

Then these are used in the call to the optimizer. In this case we assume that we want to minimize tracking error relative to `spx`.

```

> op.bhs <- trade.optimizer(prices, varian,
+   long.only=TRUE, gross.value=1.4e6, benchmark="spx",
+   existing=cur.port, lower.trade=lower.tr,
+   upper.trade=upper.tr,
+   universe.trade=c(buy.vec, sell.vec))

```

The `universe.trade` argument is restricting the assets that are traded to be only in the buy or sell list. Assets not named are constrained to not trade.

You may wish to use the `turnover` argument to constrain turnover, and other arguments.

6.4 Utility-free Optimization

The idea of utility-free optimization is that you have an ideal target portfolio and you would like to get as close to that target as possible while still obeying all of your portfolio constraints. Expected returns are not used for this and you don't have to have a variance matrix.

One of the simplest practical possibilities is a call like:

```

> op.d1 <- trade.optimizer(prices, dist.center=target.wts,
+   utility="distance", existing=current.portfolio,
+   turnover=30000, gross.value=1e6, long.only=TRUE)

```

The key parts of this call are the `dist.center` argument that takes a representation of the target portfolio, the `utility` argument, and the `turnover` argument. The call is taking advantage of the default value of `dist.style` which is `"weight"`. The `target.wts` object should be a named numeric vector that sums to 1, where the names are asset ids.

Some sort of constraint on turnover is almost surely desirable. That could be accomplished using costs, but the `turnover` argument is a much easier option.

Another constraint that is quite likely in this setting is a tracking error constraint. For this we do need a variance and the `bench.constraint` argument must be specified.

```
> op.d2 <- trade.optimizer(prices, dist.center=target.wts,
+   utility="distance", existing=current.portfolio,
+   turnover=30000, gross.value=1e6, long.only=TRUE,
+   variance=varian, bench.constraint=c(spx=0.04^2/252))
```

In this case we are specifying a maximum tracking error of 4% assuming the variance was produced with daily data—see Section 4.12 for more details.

The default style of distance (as controlled by the `dist.style` argument) is `"weight"`. In this style distance is considered to be the sum of absolute differences in the weights of the new portfolio and the target portfolio. This is often how we conceive of distance, but it has the disadvantage that the answer need not be unique. An option is to add the argument:

```
dist.style="sumsqwi"
```

This changes the distance to be the weighted sum of squared differences in weights, where the weights (for the sum of squares) are the inverse of the weights for the target portfolio. (Assets not in the target portfolio are ignored in this formulation of distance.) This form of distance may or may not be more pleasing.

Section 4.13 has more information on the arguments that control distance.

6.5 Managing Cash Flow

When you either are required to raise cash or have cash to put into the portfolio, you can use optimization to find a good trade. If there is significant cash flow, then this process has the benefit of keeping the portfolio closer to optimal than it otherwise would be—thus reducing the need for large, periodic rebalances. The examples given below assume an active portfolio—merely eliminate the `expected.return` argument if yours is passive.

Injecting Money into a Portfolio

If you want to put more money into an existing portfolio, then the main thing to say is that you don't want any selling. You probably also want to limit the number of assets to trade as well. One possible command is:

```
> op.in1 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas, turnover=newcash,
+   bench.constraint = c(spx=0.04^2/252), max.weight=.05,
+   existing=cur.port, ntrade=5, lower.trade=0)
```

The key features in this command are setting the `turnover` argument to the amount of money to be injected, and setting the `lower.trade` argument to zero.

You may also want to use the `port.size` argument to restrict the number of assets in the new portfolio. Alternatively you could set the `universe.trade` argument to be the names in the current portfolio—meaning that no additional assets will be added to the portfolio.

Extracting Money out of a Portfolio

To raise cash from the portfolio, you want the optimizer to do no buying. This is achieved with the following command:

```
> op.out1 <- trade.optimizer(prices, varian,
+   long.only=TRUE, expected.return=alphas,
+   turnover=cash.required, ntrade=5,
+   bench.constraint = c(spx=0.04^2/252), max.weight=.05,
+   existing=cur.port, upper.trade=0)
```

The pertinent parts of this command are setting the `turnover` argument to the amount of cash that is required, and setting `upper.trade` to zero.

The formulations given for both injecting and extracting money are implicitly assuming that the optimizer will want to trade at least the amount requested. That doesn't have to be true. A safer approach would be to give `turnover` an interval, perhaps something like:

```
turnover = cash.required * c(1, 1.01)
```

6.6 Asset Allocation

Asset allocation as we discuss it here means that we are interested in weights rather than prices and asset units (like shares or lots). We restrict asset allocation to long-only.

`trade.optimizer` can handle such problems.

The key step is to create a price vector that is all ones.

```
> aa.prices <- rep(1, length(asset.names))
> names(aa.prices) <- asset.names
```

You then select a gross value that is large enough to satisfy the accuracy that you desire. For example if getting the weight to the nearest basis point is good enough for you, then set the gross value to 10,000. The gross value should be given as an interval that has just one integer in it.

The command:

```
> aa.alt1 <- trade.optimizer(aa.prices, avar,
+   expected.return=aret, utility='mean-var',
+   risk.aversion=.02, long.only=TRUE,
+   gross.value=10000+c(-.5, .5))
```

performs a mean-variance optimization.

You can get the weights from the portfolio optimizer object with:

```
> valuation(aa.alt1)$weight
```

A further example of this technique is the multiple scenario analysis on page 156.

6.7 Going Farther

Tasks that you might want to undertake:

- To add trading costs to your optimization command, see Chapter 9.
- To add constraints, see Chapter 4 on page 45.
- To review common mistakes, see Section 10.1.
- To improve your solution or examine the quality of solutions you are getting, go to Chapter 12 on page 133.
- To export your solution to a file, see Section 8.3.
- To perform optimizations with multiple variances, expected returns and/or benchmarks, go to Chapter 14.

Chapter 7

Optimizing Long-Short Portfolios

This chapter discusses commands to optimize long-short portfolios.

The examples in this chapter are just sketches of what you are likely to want to do in practice. Chapter 4 discusses all the constraints that may be imposed. Some of the examples limit the number of assets that are held in the portfolio or traded. These limits are convenient, but threshold constraints (page 52) may be better to use. Chapter 9 covers trading costs—it is usually a good idea to constrain trading in some way, whether with trading costs or a turnover constraint.

7.1 Required Inputs

The `trade.optimizer` function needs a price vector. The `prices` argument needs to be a vector of prices with names that are the asset names. The assets in `prices` control the assets that are used in the problem. Other inputs—such as the variance and linear constraints—may contain additional assets (which will be ignored).

Other than `prices` there is no single argument that is required. However, there are two concepts that are required: the amount of money in the portfolio has to be constrained somehow, and there has to be something that provides a utility.

Monetary Value

The turnover or the amount of money (both long and short) in the portfolio needs to be constrained. Details can be found on page 45

Utility

At least one of the arguments:

- `variance`

- `expected.return`
- `dist.center`

must be given.

The `dist.center` argument and its mates are used to do utility-free optimization—that is, to minimize the distance to an ideal target portfolio. This is likely to be more common for long-only portfolios, but there is no reason not to do it for long-short portfolios as well. If you do want to do it, see Section 6.4 (which assumes long-only).

Assuming distance is out of the picture: If `expected.return` is not given, then the utility is the minimum variance. If a variance is not given, then the utility is the maximum expected return. If both are given, then the default utility is to maximize the information ratio.

7.2 Examples

This section presents some examples. Feel free to skip over those that don't apply to you. These examples do not include arguments for the cost of trading (see Chapter 9). Nor does it cover several of the constraints (see Chapter 4).

Maximize the Information Ratio

If you have a set of expected returns and a variance, then maximizing the information ratio is a natural thing to do—and it is the default. This can be done simply, with a command similar to:

```
> opt.ir1 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1.4e6,
+   net.value=c(-1e4, 2e4))
```

This puts 1.4 million currency units into the gross value of the portfolio, and the net value is between negative 10,000 and positive 20,000.

There are a few additional arguments that will be useful in practice. If you are building a new portfolio, then a more likely command is something like:

```
> opt.ir2 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1.4e6,
+   net.value=c(-1e4, 2e4), port.size=55, max.weight=.05)
```

This new command adds `port.size` to give a portfolio containing up to 55 assets. It also ensures that no position will be more than 5% of the gross value.

To update an existing portfolio, the command might be:

```
> opt.ir3 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1.4e6,
+   net.value=c(-1e4, 2e4), ntrade=15, max.weight=.05,
+   existing=cur.port, turnover=5e4)
```

Here both the number of assets in the trade and the turnover are limited, and of course the existing portfolio is specified.

Maximize Return with a Bound on the Variance

If you have a target for the maximum variance that you want, then you can add a constraint on the variance. Suppose that volatility should be no more than 4%. We need to transform the 4% into a variance, and (possibly) go from annual to the frequency of the variance. If the variance is on a daily frequency (and created from returns in decimal as opposed to percent), then the command would look like:

```
> opt.vc1 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1.4e6,
+   net.value=c(-1e4, 2e4), ntrade=55, max.weight=.05,
+   var.constraint=.04^2/252)
```

This is not doing quite what is advertised in the headline, but may actually be more appropriate. This is really maximizing the information ratio with a constraint on the variance. If the best information ratio has a variance that is smaller than the bound, then it will be the best information ratio that is chosen.

If you truly want the maximum return given the variance bound, then you need to say so explicitly:

```
> opt.vc2 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1.4e6,
+   net.value=c(-1e4, 2e4), ntrade=55, max.weight=.05,
+   var.constraint=.04^2/252, utility='maximum return')
```

Minimize Variance Given a Long List and a Short List

Suppose you don't have specific predictions for the returns of the assets, but you do have a set of assets that you predict will go up and a set you think will go down. That is, you have a buy-hold-sell list. In this case you can force only buys from the long list, sells from the short list, and minimize the variance of the portfolio.

If `long.names` and `short.names` are vectors containing the asset names that are predicted to go up and down, then the computation can proceed along the lines of:

```
> long.zero <- rep(0, length(long.names))
> names(long.zero) <- long.names
> short.zero <- rep(0, length(short.names))
> names(short.zero) <- short.names
```

S code note

These commands create two new vectors containing all zeros and with names that are the long assets for one of them, and names that are the short assets for the other. The `rep` function replicates its first argument.

These new vectors are then used as the `lower.trade` and `upper.trade` arguments:

```

> opt.list1 <- trade.optimizer(prices, varian,
+   gross.value=1.4e6, net.value=c(-1e4, 2e4),
+   lower.trade=long.zero, upper.trade=short.zero,
+   universe.trade=c(long.names, short.names),
+   ntrade=50)

```

The `lower.trade` and `upper.trade` arguments to `trade.optimizer` are limits on the number of units (e.g., lots or shares) that can be traded. Each asset can be given a different limit. If what is passed in is a vector with names, then any number of assets may be represented (in any order). So using `long.zero` as the `lower.trade` argument says that the assets named by `long.zero` can not have negative trades—that is, can only be bought. The same logic applies to the `upper.trade` argument with `short.zero`.

Another important piece of the command is setting `universe.trade` so that assets that are in neither `long.names` nor `short.names` won't trade at all.

This example is a simple case—the lower and upper bounds can be more complicated when there is an existing portfolio. It is possible that using the `positions` argument (page 54) would be easier in such a case.

If, for example, the existing portfolio has a long position in an asset that is on the short list, you may wish to force a sell in the asset that is at least as large as the existing position. The commands above will not do that—you need to resort to either the `forced.trade` argument or the `positions` argument.

```

> opt.list2 <- trade.optimizer(prices, varian,
+   gross.val=1.4e6, net.val=c(-1e4, 2e4),
+   lower.trade=long.zero, upper.trade=short.zero,
+   universe.trade=c(long.names, short.names),
+   ntrade=50, forced.trade=c(ABC=-25))

```

In this example we are forcing a sell of at least 25 units of asset ABC.

Mean-Variance Optimization

To do mean-variance optimization, you need to give both a variance and expected returns and set:

```
utility = "mean-variance"
```

You will want to specify the risk aversion as well. More details of the risk aversion parameter are on page 140.

```

> opt.mv1 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1.4e6,
+   net.value=c(-1e4, 2e4), ntrade=55, max.weight=.05,
+   utility='mean-var', risk.aversion=3.7)

```

Alternatively, you can do mean-volatility optimization by using:

```
utility = "mean-volatility"
```


7.3 Managing Cash Flow

The optimizer can be used to select trades to take money out of the portfolio, or put money into it. The optimizer makes this process quite painless, even when you want a fair amount of control over the trade.

The examples given below assume that the cash management is via the gross value. An alternative would be to control the longs and shorts directly. You would use the `long.value` and the `short.value` arguments for this more specific control.

Injecting Money into a Portfolio

If you want to perform optimization with a constraint on the variance (see page 103), then a command that will do this while only increasing existing positions is:

```
> curgross <- valuation(cur.port, prices)$total["gross"]
>
> opt.in1 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=curgross + newcash,
+   net.value=c(-1e4, 2e4), ntrade=15, max.weight=.05,
+   var.constraint=.04^2/252, existing=cur.port,
+   turnover=newcash, universe.trade=names(cur.port))
```

The key points here are setting the gross value to the desired new amount, and forcing the turnover to equal the amount by which the gross is increased—this means that no positions can be reduced. Setting the trade universe to be the assets that are in the current portfolio ensures that no new positions are opened. If it is okay to open new positions, then do not give the `universe.trade` argument. You can also increase `turnover` any amount that you want in order to do rebalancing as well as managing cash.

Extracting Money out of a Portfolio

Taking money out of the portfolio is similar to putting it in, a suitable command might be:

```
> curgross <- valuation(cur.port, prices)$total["gross"]
>
> opt.out1 <- trade.optimizer(prices, varian,
+   expected.return=alphas, net.value=c(-1e4, 2e4),
+   gross.value=curgross - cash.required, ntrade=15,
+   turnover=cash.required, max.weight=.05,
+   var.constraint=.04^2/252, existing=cur.port,
+   universe.trade=names(cur.port))
```

The comments about the `universe.trade` and `turnover` arguments made about injecting money also apply here.

The formulations given for both injecting and extracting money are implicitly assuming that the optimizer will want to trade at least the amount requested. That doesn't have to be true. A safer approach would be to give `turnover` an interval, perhaps something like:

```
turnover = cash.required * c(1, 1.01)
```

7.4 Money Constraints

If you are looking for a control in the optimizer for leverage, there isn't one. The amount of leverage applied is determined by the fund manager—leverage is external to the optimization. The decision on leverage comes first, which determines the gross value and the net value—quantities that the optimizer does need to know.

The constraints on the long and short values are not just redundant to the gross and net value constraints. Figure 7.1 shows that if we plot the gross and net value constraints, then these form a rectangle that is diagonal to the long value and short value axes. The constraints displayed in the figure are the gross value between 10 and 12, and the net value between -1 and 2. Constraints on the long value and the short value will be vertical and horizontal sections, so adding these change the shape of the allowable region.

7.5 Real-Time Monitoring

You may want to perform real-time optimization. The prices will change from minute to minute, and if your expected returns change continuously as well, then constantly updating the optimization could be valuable. Such a process is very easy to set up. Here is an example:


```
> repeat{
+   prices <- update.prices() # user-defined function
+   alphas <- update.alphas() # user-defined function
+   cur.port <- update.portfolio() # user-defined function
+   cur.opt <- trade.optimizer(prices, varian,
+     expected.return=alphas, turnover=5e4, ntrade=5,
+     existing=cur.port)
+   if(length(cur.opt$trade)) {
+     deport(cur.opt, what="trade")
+   }
+ }
```

S code note

In S, `repeat` performs an infinite loop. Generally the body of a repeat loop will contain a test of when the loop should be exited. In this case there is no such test—the optimizations will continue to be performed until you explicitly interrupt the loop. It would be easy enough to use the `date` function to check the time of day and exit the loop after the close of trading.

This example has three functions that you would need to write. These functions update the prices, the expected returns, and the positions in your

Figure 7.1: Constraints on gross, net, long and short values.



portfolio. If you are willing to assume that the suggested trades always happen, then you could merely have:

```
cur.port <- cur.opt$new.portfolio
```

The technique here might be characterized as looking for the next small thing—the size of the trade and the maximum number of assets to trade should probably be small.

7.6 Going Farther

Tasks that you might want to undertake:

- To add trading costs to your optimization command, see Chapter 9 on page 115.
- To add constraints, see Chapter 4 on page 45.
- To review common mistakes, see Section 10.1 on page 121.
- To improve your solution or examine the quality of solutions you are getting, go to Chapter 12 on page 133.
- To export your solution to a file, see Section 8.3 on page 113.
- To perform optimizations with multiple variances and/or multiple expected return vectors, see Chapter 14.
- If you want a benchmark with a long-short portfolio, see Section 11.2.

Chapter 8

General Use

The core functionality of Portfolio Probe is generating random portfolios and optimizing trades. More detailed description of these activities can be found in other chapters. This chapter focuses on peripheral tasks. In particular there are hints about using the S language.

8.1 Setting Up Data

Before computations can be performed, there needs to be data in place on which to do the computations. Data that we may need include prices, existing positions in the portfolio, expected returns, variances, membership of assets in various groups such as countries and sectors, and so on.

Prices and Other Imports

You will probably need to import prices into S. You may need to get other data into S as well. A typical command to create a price vector is:

```
> prices <- drop(as.matrix(read.table("prices.txt",
+   sep="\t")))
```

or

```
> prices <- drop(as.matrix(read.table("prices.csv",
+   sep=",")))
```

depending on whether your file is tab-separated or comma-separated. (Other file formats can be handled as well.)

S code note

These commands are assuming that there are two items on each line of the file—an asset name and a number which is the price of a given quantity of the asset. The `read.table` function returns a type of S object called a *data frame*. In S data frames and matrices are both rectangles of entries. A matrix must have all of its entries the same type—for example all numbers or all character

strings. A data frame has the freedom to have different types of entries in different columns—one column could be numeric and another categorical.

In this example `read.table` creates a data frame that has one column, which is numeric, and the asset names are used as the row names. This is then converted from a data frame into a matrix. Finally the `drop` function changes the object from being a one-column matrix into an ordinary vector. The reason that there is the intermediate step of converting to a matrix is that the row names are not preserved when a data frame is dropped to a vector. (There is a reason for this—you can just think of it as an unfortunate fact.)

Other data will be imported in a similar fashion. Using `as.matrix` is fairly likely, `drop` is only appropriate when there is a single column (or row) of data and a simple vector is desired.

Variance Matrix

A variance matrix is often used in Portfolio Probe. If your problem is an asset allocation with at most a few dozen assets, then a sample variance matrix as computed by the `var` or `cov.wt` S functions should do (though there may be better methods). For a large number of assets, a factor model or a shrinkage model will probably be more appropriate.

You can use any variance matrix that you want. It can come from a commercial source or you can compute it yourself.

If you need to compute the variance matrix yourself, there are functions in the `BurStFin` package that will do that.

If you are using R on Windows, then you can get `BurStFin` with:

```
> install.packages("BurStFin",
+   repos="http://www.burns-stat.com/R")
```

You can get the source with:

```
> download.packages("BurStFin", 'some_directory',
+   type="source", repos="http://www.burns-stat.com/R")
```

Once the package is installed for your version of R, you can do:

```
> library(BurStFin)
```

in each session in which you want to use its functionality (this command might be slightly more complicated if you are using S-PLUS).

The `factor.model.stat` function in `BurStFin` creates a variance matrix using a factor model. The command to get a variance matrix can be as simple as:

```
> retmat <- diff(log(price.history))
> varian <- factor.model.stat(retmat)
```

Alternatively, a variance that shrinks towards the equal correlation matrix can be computed with:

```
> varian <- var.shrink.eqcor(retmat)
```

S code note

Here `price.history` is a matrix of asset prices with the columns corresponding to the assets and the rows corresponding to times (the most recent time last). The calculation produces a matrix of returns which we call `retmat` in this case. (The vector that is used as the `prices` argument could be the final row of the price history matrix.)

The `log` function returns an object that looks like its input, but each element of the output is the natural logarithm of the corresponding element of the input.

The `diff` function when given a matrix performs differences down each column. The first row of the result is the second row of the input minus the first row of the input; the second row of the result is the third row of the input minus the second row of the input; *et cetera*.

note

The `retmat` object holds log returns (also known as continuously compounded returns). You may hear strong opinions on whether simple returns or log returns are preferred. In mean-variance optimization neither is entirely correct, however the results are likely to be quite similar—especially if daily or higher frequency returns are used.

caution

Do not use daily returns when your data include assets from markets that close at different times—global data in particular. The asynchrony of the returns means that the true correlations are higher than those that will be estimated from the data. Thus optimizations will be distorted. Weekly is the highest frequency that should be used with data that are substantially asynchronous. (Note though that techniques do exist to adjust for asynchrony—see [Burns et al., 1998].)

caution

There is the possibility of confusion with the word “factor”. In S an object that is of type “factor” is something that is categorical. For example, a vector of the sector or country of assets will (at some point at least) be a factor when used as a linear constraint. A factor model does not contain any factors in this sense.

A large amount of the effort in `factor.model.stat` is dealing with missing values. The default settings are reasonable for long-only optimizations. Possibly a more appropriate matrix for benchmark-relative analyses and long-short optimization would be:

```
> varian <- factor.model.stat(retmat, zero=TRUE)
```

For assets with missing values this biases correlations toward zero rather than towards the average correlation as the previous command does.

advanced

The treatment of missing values can have a material impact on the optimization, and it can be beneficial to specialize the estimation to your particular application. An easy way to do this is to return the factor model representation:

```
> varfac <- factor.model.stat(retmat, out="factor")
```

When the `output` argument is set to `"factor"`, the result is a list with components named `loadings`, `uniquenesses` and `sdev`. You can modify these to correspond to an approach appropriate for you. Then you can create the variance matrix from the object:

```
> varian <- fitted(varfac)
```

This works because the result of `factor.model.stat` when it returns the factor representation—`varfac` in this case—has a class, and that class has a method for the `fitted` generic function. The result of this function is the variance matrix that is represented by the factor model.

Adding a Benchmark to the Variance

Perhaps you have a variance matrix of the individual assets, but the benchmark is not included in the variance. You can add the benchmark to the variance matrix if you have the weights for the constituents of the benchmark and all of the constituents are in your variance matrix. Merely use the `var.add.benchmark` function from `BurStFin`:

```
> varian <- var.add.benchmark(varian.assets, spx.weights,
+ "spx")
```

This function takes three arguments—a variance matrix, a vector of weights (that should sum to 1), and a character string giving the name for the benchmark. The weight vector needs to have names which are (some of) the asset names on the variance matrix.

The computation performed by this function is the preferred method of introducing a benchmark into a variance matrix for optimization. See [Burns, 2003a] for a study on approaches to incorporating benchmarks into the variance matrix for optimization.

8.2 The Random Generation or Optimization

The optimization process itself is covered in other chapters. The basics are covered in chapters that depend on what you are doing:

- Random portfolio generation: Chapter 3.
- Long-only portfolio optimization (active, passive or utility-free): Chapter 6.
- Long-short portfolio optimization: Chapter 7.

Trading costs are discussed in Chapter 9, while Chapter 4 covers constraints.

Chapters 10 and 14 have additional information—Section 10.1 on page 121 is particularly recommended.

caution

If in optimization a penalty is imposed because one or more constraints are broken, then the optimizer will produce a warning. Do not ignore the warning. It could well mean that there is a mistake in the specification.

You can re-run the optimization—perhaps with an increased number of iterations—to see if the optimizer finds a solution that does not break any constraints. If that doesn't work, you should probably investigate if you have imposed constraints that don't allow a solution. Section 10.1 on page 121 lists some common mistakes.

On the other hand, this method of treating constraints allows backtests to proceed even when the constraints that are imposed are infeasible for some of the time periods. The answers that break the constraints, while not optimal, may be (at least sometimes) reasonable.

8.3 Post-Optimization

This section is specific to optimization. If you are generating random portfolios, then the analogous discussion is Section 3.3.

Once you have performed an optimization, there are two tasks remaining—see what the trade is, and export the data so that the trade can be implemented.

Explore the Trade

Information on the optimization can be viewed using the `summary` function:

```
> opti <- trade.optimizer(prices, varian, long.only=TRUE,
+   gross.val=1e6)
> summary(opti)
```

The `summary` function provides information on:

- the utility, costs, and penalties for broken constraints
- opening and closing positions, etc.
- valuations of the portfolio and trade
- realization of linear (and count) and distance constraints

You can get some more specific information by printing the optimal portfolio object (which can be done by just typing its name):

```
> opti
```

Alternatively, you could look at an individual component of the object, such as the trade:

```
> opti$trade
```

The prices that `summary` uses to value the portfolio are, by default, the prices used in the optimization. You can give `summary` different prices:

```
> summary(opti, price=new.prices)
```

You can also get the valuation of the individual assets as well as overall valuation using the `valuation` function:

```
> valuation(opti)
> valuation(opti, trade=TRUE)
```

The default behavior of `valuation` when given a portfolio object is to give information on the portfolio. The trade can be examined instead by setting the `trade` argument to `TRUE`.

When `valuation` is given a vector, the prices must be given:

```
> valuation(opti$trade, price=new.prices)
```

Prices are optional when portfolio objects are given—the prices used during the optimization are used if none are specified.

Export the Trade

The other task to be performed—once you are satisfied with the optimization—is to place the information elsewhere. The `deport` function does this (“export” is a logical name for this functionality, but that has a different meaning in R). The simplest use of this function is just to give the name of the portfolio object:

```
> deport(opti)
[1] "opti.csv"
```

The return value of the function is a character string giving the name of the file that was created. If the optimization was performed in terms of lots but you want the file to reflect number of shares, the `mult` argument should be given. If, for instance, all of the lot sizes are 100, then you would do:

```
> deport(opti, mult=100)
[1] "opti.csv"
```

When there are different lot sizes, you should give a vector containing the lot size of each asset:

```
> deport(opti, mult=lotsizes)
[1] "opti.csv"
```

If you want the file to be in monetary units, then the `mult` argument should be prices:

```
> deport(opti, mult=prices, to="txt")
[1] "opti.txt"
```

This command writes a tab-separated file (which in this case is called `opti.txt`). The `what` argument can be given to control what information is put into the file—see the help file for details.

Chapter 9

Trading Costs

This chapter covers trading costs. It is possible to constrain costs, so costs can be used when generating random portfolios. This is only likely when imitating a series of trades with random trades.

Costs are a tricky part of optimization, but are extremely important.

9.1 Background

If there were any justification for the attitude that optimization is hard, it should be on account of trading costs. However, the problem of trading costs is fundamental to fund management—you have to confront the issue whether or not you use an optimizer.

There are numerous sources of trading costs—commissions, bid-ask spread, and market impact are commonly considered. Less obvious costs can be such things as capital gains taxes if a position is sold, and the cost of imperfectly known expected returns. Costs need not be symmetric for buying and selling. You may, for example, want the cost of increasing positions to be more than the cost of decreasing them, because of increased liquidity risk and/or because of the implicit cost of needing to close the position again.

9.2 Specifying Costs

Costs are specified similarly in `trade.optimizer` and `random.portfolio`. When generating random portfolios, costs will be ignored unless `limit.cost` is given (or the utility is constrained).

The specification of costs is very flexible while simple costs can be given easily. To allow for asymmetry in costs, there are four cost arguments:

- `long.buy.cost`
- `long.sell.cost`
- `short.buy.cost`
- `short.sell.cost`

This allows differences in the costs for long positions versus short positions, as well as differences for selling and buying.

If there is no difference between buying and selling, and between long and short positions, then give only the `long.buy.cost` argument.

caution

If you give one of the cost arguments other than `long.buy.cost` but no others, then the named situation will be the only case where there is cost. That seems unlikely to be what you want.

Minimally, costs must be given for all of the assets that are allowed to trade in the problem (if costs are given at all). Costs may be given for any number of assets as long as all the tradeable assets are included.

Linear Costs

Linear costs are given with a vector (or a one-column matrix) of costs. For example, to have costs of 30 basis points for all trades, the command would be:

```
> opt1 <- trade.optimizer(prices, ...,
+   long.buy.cost = prices * .003)
```

To have costs of 50 basis points for short sales and costs of 30 basis points for all other trades, the command is:

```
> opt2 <- trade.optimizer(prices, ...,
+   long.buy.cost = prices * .003,
+   short.sell.cost = prices * .005)
```

The “...” in these commands and subsequent ones refers to additional arguments that you want in the optimization.

The advantage of linear costs is that they are easy to think about and easy to write down. Linear costs are undoubtedly more popular than they otherwise would be because computational engines have generally been unavailable for more realistic models of cost.

Nonlinear Costs

Costs are unlikely to really be linear. It can be quite useful to have non-integer exponents in the trading cost functions. This is allowed via the `cost.par` argument. `cost.par` is a vector giving the exponents for the cost function.

Consider:

```
> cost.poly
      [,1] [,2] [,3] [,4]
ABC 0.00 0.01 0.02 0.00
DEF 0.03 0.02 -0.01 0.01
> trade.optimizer( ..., long.buy.cost=cost.poly)
> trade.optimizer( ..., long.buy.cost=cost.poly,
+   cost.par=0:3)
```

These two calls to `trade.optimizer` are equivalent (though there would be a slight computational efficiency advantage to the first one). Polynomial costs are discussed below.

A more telling use of `cost.par` is:

```
> trade.optimizer( ..., long.buy.cost=cost.poly,
+   cost.par=c(0, 1, 1.5, 2.34))
```

In this case if ABC trades 30 lots and DEF trades -25 lots, then the cost for ABC will be:

$$0 + .01(30) + .02(30^{1.5}) = 3.586$$

The cost for DEF will be:

$$.03 + .02(25) - .01(25^{1.5}) + .01(25^{2.34}) = 17.952$$

When `cost.par` is used, there is a restriction that all of the cost matrices must have the same number of columns as `cost.par`. In actuality there is no loss of generality as columns of zeros can be added to matrices as required, and `cost.par` can be the union of all of the desired exponents.

We'll see in the next section that in a long-short portfolio when a trade makes an asset go from long to short (or vice versa) that there is a choice of having the intercept counted in both instances or just the first. There is no such choice when using `cost.par`—it doesn't know about zero being the intercept and the full cost is always used.

Polynomial Costs

This subsection describes creating polynomial costs. Using the more general `cost.par` argument is usually preferred, so feel free to skip this subsection.

Polynomial costs are given by a matrix. The order of the polynomial is one less than the number of columns in the matrix. The first column contains the intercept, the second column has the linear coefficients, the third column has the quadratic coefficients, etc.

caution

Do not forget the intercept.

```
> WRONG <- cbind(linear.coefs, quad.coefs)
> okay.costs <- cbind(0, linear.coefs, quad.coefs)
```

Because it is easy to forget the intercept, a warning is issued if there are non-zero values in any intercept. If you really do have non-zero elements in cost intercepts and you don't want to see the warning about it, then you can set the `do.warn` argument to `c(cost.intercept.nonzero = FALSE)`. Only the first part of the name needs to be given—enough needs to be given so that it is unambiguous relative to the menu of names that can be given `do.warn`.

Consider this example matrix given as the costs:

```
> cost.poly
      [,1] [,2] [,3] [,4]
ABC 0.00 0.01 0.02 0.00
DEF 0.03 0.02 -0.01 0.01
```

If ABC trades 30 lots and DEF trades -25 lots, then the cost for ABC will be:

$$0 + .01(30) + .02(30^2) = 18.3$$

The cost for DEF will be:

$$.03 + .02(25) - .01(25^2) + .01(25^3) = 150.53$$

Note that the matrix needs to have enough columns to encompass the highest order polynomial—lower order polynomials will have zeros for their higher order coefficients.

caution

You should make sure that each polynomial that you give makes sense throughout the allowable trading range of the asset. Negative costs, for instance, are often not reasonable. (Though there are special circumstances—such as crystallizing a loss for tax purposes—where negative costs are appropriate.)

If your costs are just intercepts—there is a fixed cost for trading any of an asset, then you need to give a two-column matrix with the second column containing all zeros:

```
> intercept.only <- cbind(intercepts, 0)
```

This is most sensible if the intercepts are not all equal—otherwise it is almost the same as constraining the number of assets to be traded.

If the portfolio is long-short, then there can be more than one type of trading cost for a single asset. For example, an asset that is long in the existing portfolio can be sold heavily enough that it becomes short. There is then the issue of whether or not both intercepts should be counted in the trade. The `doubleconst` control parameter determines which is done. If `doubleconst` is `FALSE` (the default), then only the first intercept is used. If `doubleconst` is `TRUE`, then both intercepts are counted.

9.3 Power Laws

[Grinold and Kahn, 2000] present an argument that the trading cost per unit of an asset should be commission plus half the bid-ask spread plus some constant times the square root of the quantity: the amount to be traded divided by the average daily volume. The last term involving the square root comes from an inventory model. They suggest that the constant should be on the order of the daily volatility of the asset.

Sample code implementing this model might look like:

```
> cost.mat <- cbind(commission + spread[names(prices)] / 2,
+   volatility[names(prices)] / sqrt(ave.daily.volume[
```

```

+   names(prices]))
>
> trade.optimizer(..., long.buy.cost=cost.mat,
+   cost.par=c(1, 1.5))

```

The first command creates a two-column matrix. The first column relates to the items that are linear (commission and spread), the second column contains the cost involving the square root of the trade volume. When this is used in an optimization, the exponents for these columns are 1 and 1.5, respectively.

The second coefficient is 1.5 (where naively you may expect it to be 0.5) because the square root rule says that the cost *per share* depends on the square root of the number of shares traded. (If you think in terms of calculus, we are taking the integral of the cost function.)

[Almgren et al., 2005] empirically found a power of 0.6 to fit a set of trades on large US stocks. Since a square root is a power of 0.5, their findings suggest that trading costs grow slightly faster as the size of the trade increases. If you were to use this power law, then the `cost.par` in the example above would be 1 and 1.6 instead of 1 and 1.5.

9.4 On Scaling Costs Relative to Utility

There are (at least) three things we want in the costs that we declare:

- The cost of a larger trade in an asset is appropriate relative to a smaller trade of that asset.
- The cost of trading one asset is well-scaled relative to the cost of trading the other assets.
- The costs are appropriately scaled in light of the utility that is used.

The proper scaling of the costs relative to the utility—whether that is expected returns (and variance), or distance—is the point where the science gets a bit fuzzy. This is also a key issue controlling the merit of the optimization. If the costs given the optimizer are too small, then excessive trading eats up performance. If the given costs are too large, then performance is hurt via lost opportunity. Perfection is not necessary for the optimization to be of value, but attempting perfection is probably not such a bad mission.

The added utility of a trade needs to compensate for the transaction cost. Thus the costs should correspond to the expected returns over the expected holding period.

To be more specific, suppose that we think it will cost 50 basis points to trade a certain asset. If we are using daily data and our expected return for this asset is 10 basis points, then we will have to have a holding period of 5 days in order to break even. To amortize the cost, it (that is, the 50 basis points) should be divided by the number of days that we expect to hold this asset. If we expect to hold it for 2 days, the optimizer would see a cost of 25 basis points. If we expect to hold it 10 days, the optimizer would see a cost of 5 basis points. (The statement of this example might seem to imply that volatility is of no concern—it is of concern.)

Let's revisit the example from Section 9.3:

```

> cost.mat <- cbind(commission[names(prices)] +
+   spread[names(prices)] / 2,
+   volatility[names(prices)] / sqrt(ave.daily.volume[
+   names(prices)]))
>
> cost.amort <- cost.mat / expected.holding[names(prices)]
>
> trade.optimizer(..., long.buy.cost=cost.amort,
+   cost.par=c(1, 1.5))

```

This adds a command which divides (what might be called) the raw costs by how long we expect to hold each asset. The result is the amortized costs. The holding period should be in the same time units as the expected returns. If we have daily returns, then an asset that we expect to hold 25 days should have 25 as its expected holding period. However that same asset should have an expected holding period of about 0.1 when we are using annualized returns.

That the expected returns are noisy suggests that the costs should be increased from this “rational” value to reduce trading.

There is another sense of scaling costs. By default `trade.optimizer` sums the costs from all of the trading and then divides by the gross value of the portfolio. This is the natural thing to do as the expected returns and the variance are also scaled by the gross value when the utility is computed. However, you do have the option to scale by the trade value or not to scale the costs at all—this is controlled by the `scale.cost` argument.

9.5 Costs Due to Taxes

Taxes present a cost structure that is entirely different from trading costs resulting from the market. In particular the cost will be different depending on whether you are buying or selling. More detail on costs due to taxes can be found in [diBartolomeo, 2003] and its references.

Several optimizers implement piecewise linear cost functions. While this choice is likely to be due to mathematical tractability, costs (and benefits) from taxes really are piecewise linear in some circumstances. The tax liability often depends on the length of time an asset has been held, and there may be several holding periods for a particular asset. Portfolio Probe does not have piecewise linear costs, but you can approximate them with nonlinear costs.

Taxes also have a sort of time decay similar to options. Their effect changes as the end of the tax period gets closer.

9.6 Going Farther

Additional steps might be:

- To look for trouble spots with Chapter 10 on the facing page

Chapter 10

Practicalities and Troubleshooting

The best question to ask after a computation has been performed is: **Does it make sense?**

The purpose of this chapter is to help with that question—to give hints about how it might not make sense, and to give possible solutions when it doesn't.

Though there will be a tendency to come to this chapter only when you think something is wrong, the best time to come to it is when you think everything is right.

While some of the problems listed in this chapter are specific to optimization, there are many that will affect the generation of random portfolios also.

The sections of this chapter are:

1. Easy ways to be wrong.
2. Suppressing warnings (page 124).
3. Cheatsheets (page 126).
4. Troubleshooting (page 127).
5. S Language Problems and Solutions (page 129).

10.1 Easy Ways to Be Wrong

The phrase “garbage in, garbage out” has gone slightly out of fashion, but it still has currency. What follows are a few possibilities for garbage.

Data Mangling

The prices are not all in the same currency.

It doesn't matter which currency is used, but there needs to be only one.

There are prices for the wrong quantity.

For example, the optimization is thought of in terms of lots, but the prices for some or all of the assets are for shares.

There are stock splits, rights issues, etc. that have not been accounted for.

Stock splits and rights issues that are not caught can mean that your existing holdings are wrong in the optimization.

A missed split in the price history degrades variance estimates that are based on returns created from the price history.

The variance is of prices, not returns.

The variance matrix of the returns of the assets is needed. If the variance is of the prices, the results will be seriously wrong. (The functions in `BurStFin` that create variance matrices from data give a warning if prices are used.)

The variance and/or expected returns are not properly scaled.

The scaling that is chosen does not matter, but it needs to be consistent. The (time) scales of the variance and the expected returns should match—for example they can both be annualized, or both be daily. You also need such things as benchmark constraints to be scaled to match the variance. Costs need to account for the time scale of the expected returns, and whether or not the expected returns are in percent.

The variance matrix is singular.

While the Portfolio Probe optimization algorithm will happily handle singular matrices, the results need not be useful. If there are more assets than time points in the returns data, then a factor model or shrinkage estimator rather than a sample variance should be used. This problem is certainly more serious for long-short portfolios, but is problematic for long-only portfolios as well.

If the true variance really is singular—for example cash is one of the assets (see page 156)—then you should get good results. The problem is if the variance matrix says that there are portfolios that are essentially riskless when they are not.

The variance matrix is not symmetric.

The optimizers inherently assume that the variance is symmetric. They will get confused if this is not the case. If what you are using as the variance is not symmetric, then you can get a version that is with:

```
> varian.sym <- (varian + t(varian)) / 2
```

There is asynchrony in the data.

Global data should not be used at a higher frequency than weekly unless asynchrony adjustments are made. In addition to different opening hours, asynchrony can be caused by assets that do not trade often. The illiquid assets will have stale prices. A model that adjusts for asynchrony due to different opening hours is presented in [Burns et al., 1998].

The expected return for a benchmark is not equal to the weighted average of the expected returns of its constituents.

This problem is pointed out in [Michaud, 1998]. Minor violations are unlikely to have much affect, but differences do distort the optimization.

The signs of covariances in the variance matrix are reversed for assets that are intended to be short.

While this could produce valid results if properly done, it is confusing and totally unnecessary. Portfolio Probe was developed with long-short portfolios in mind—bizarre manipulations can be eliminated.

Input Mangling

While many mistakes are caught by the software, it is quite possible to confuse the inputs. One check is to see if the optimizer is using the objective that you intend. You can do this by:

```
> opt.1$objective.utility  
[1] "mean-variance, risk aversion: 0.02"
```

The name of an argument is inadvertently dropped.

For example, you add a starting solution to a command, but forget to name it `start.sol` so the optimizer thinks it is a different argument. Most arguments are checked for validity, but in some cases the checks may not be good enough.

You confuse the meaning of “start.sol” with “existing”.

The `start.sol` argument should be one or more solutions that are used as starting values for the trade in optimization. The current portfolio should be given in the `existing` argument.

You confuse the meaning of the “lower.trade” or “upper.trade” argument.

If you habitually abbreviate the `lower.trade` argument to `lower`, you may begin to think of the bound as being on the final portfolio rather than on the trade.

Bounds for linear constraints are not properly scaled.

Values in `lin.bounds` need to match the corresponding value in `lin.style`. If the bounds are in values but the style is in weight, then there probably isn't a real constraint. Alternatively if the bounds are in weights and the style is in value, then the constraint is likely to be impossible.

You do not give the `lin.abs` or `lin.trade` argument when needed, or reverse the meaning.

If the bounds are set for absolute constraints but `lin.abs` is not given, then the constraints may be inconsistent. Perhaps even worse, they may not be inconsistent and give an answer that on the surface seems okay.

You give just one trading cost argument, but it is the wrong argument.

If you want costs to be the same whether the position is long or short, and whether you are buying or selling, you need to give the `long.buy.cost` argument. If you give a different argument, then the costs correspond to the named situation only and the costs in other situations are zero. There will be a warning if you do this.

10.2 Suppressing Warning Messages

While warning messages help avoid mistakes, they can be annoying (and counter-productive) if they warn about something that you know you are doing. The `do.warn` argument of `random.portfolio` and `trade.optimizer` can be used to suppress most warnings. Here is an example of its use:

```
do.warn = c(utility.switch=FALSE, value.range=FALSE)
```

Note that abbreviation of the names is allowed as long as the abbreviation is unique among the choices.

Some warnings are never suppressed—these are more likely to be mistakes. You should reformulate the command to avoid such warnings whether or not the original is a mistake.

The default is `TRUE` for all entries except `converged` (which defaults to `FALSE` since non-convergence of the optimization is seldom of concern).

Here are the possible entries for `do.warn`:

- **benchmark.long.short**: A benchmark is used in a long-short portfolio. While this can be correct, it often is not. See Section 11.2.
- **bounds.missing**: There are rows missing from the `lin.bounds` object that should logically be there. The missing bounds are taken to be infinite. This is almost certainly caused by something going wrong, or at least sloppy use.

- **converged:** (optimization only) Convergence is declared if more than `fail.iter` consecutive iterations fail to improve the solution. If the maximum number of iterations is performed before this condition occurs, then there is no convergence for the run. When `stringency` is positive, the convergence is a little more complicated.
- **cost.intercept.nonzero:** One or more of the cost arguments have non-zero values for the intercept. (This only applies when polynomial costs are given.) This may well be correct, but it is easy to forget the intercept column in the cost matrix.
- **dist.prices:** One or more of the assets that are in `prices` are not in (a component of) `dist.prices` and hence are assumed to be zero.
- **dist.style:** Since it is easy to make a mistake between giving shares or values to the `dist.center` argument, a check is made to see if what is stated in `dist.style` fits into the gross value range that is given or inferred.
- **dist.zero:** One or more of the assets in (a component of) `dist.center` have a non-zero value but a price of zero in `dist.prices`.
- **extraneous.assets:** One or more objects contain data on assets that are not in `prices` and hence not of use in the computation.
- **ignore.max.weight:** One or more assets in the existing portfolio break their maximum weight constraint; and either maximum weights are not forced to be obeyed, or the trade can not be forced to be large enough. See page 49.
- **infinite.bounds:** All bounds in `lin.bounds` are infinite, meaning there are no bounds at all.
- **max.weight.restrictive:** The vector of maximum weights is very restrictive. That is, the sum of (a subset of) the maximum weights is only slightly more than 1.
- **neg.dest.wt:** (optimization only) There is at least one negative destination weight in the utility table. This is almost surely wrong unless the problem is outside typical portfolio optimization.
- **neg.risk.aversion:** (optimization only) There is at least one negative risk aversion value. Risk aversion is traditionally thought to be non-negative, but you could have a special situation.
- **no.asset.names:** One or more objects do not have asset names, and hence require that they be in the same order as the assets in `prices`. Some objects are allowed to not have asset names in order to facilitate quick experiments—asset names are always recommended for production situations.
- **noninteger.forced:** There is a forced trade that is constrained such that the trade must be non-integer. This could be something you want to do, but can mean that something is wrong with the specification of constraints.

- `nonzero.penalty`: (optimization only) The penalty is non-zero, meaning that not all constraints are satisfied.
- `notrade`: The formulation does not allow any trading. If you really want no trading, then the recommended approach is to set `ntrade = 0`.
- `novariance.optim`: (optimization only) An optimization (other than with a distance) is requested but no variance is given. This can be a reasonable operation for some long-short situations, but is generally suspicious.
- `penalty.size`: (optimization only) The risk aversion parameter is large relative to the values in `penalty.constraint`. In this case the optimizer may have problems getting the constraints to be obeyed.
- `positions.names`: The asset names in the `positions` argument are not the same as the universe of assets in the problem.
- `random.start`: (random portfolios only) The `start.sol` argument is given, which is ignored in `random.portfolio`. `start.sol` is sometimes confused with `existing` (the current portfolio).
- `randport.failure`: (random portfolios only) Fewer than the requested number of random portfolios were generated.
- `start.noexist`: The `start.sol` argument is given, but `existing` is not. This is sometimes what is wanted, but it might indicate that the meaning of `start.sol` is confused with `existing`.
- `turnover.max`: The maximum turnover is too large to have an effect. If you want to specify a minimum turnover but no maximum turnover, then this warning is avoided by setting the maximum to `Inf`.
- `utility.switch`: (optimization only) The utility (either given or default) is switched to another utility.
- `value.range`: One or more of the ranges for monetary value (`gross.value` etc.) is narrow relative to the prices.
- `variance.list`: A compact variance is given, which means that there is no way to check that the assets are the same (and in the same order) as the assets in `prices`.
- `zero.iterations`: The `iterations.max` argument can be set to zero, but this is not the same as doing no optimization. To get the `start.sol` trade as the final trade, set `funeval.max` to zero or one.

10.3 Cheatsheets

Implied Ranges

Table 10.1 lists the true meaning of a single number for those arguments that represent a range of values. Note that `close.number` is the outlier in that its single value means exactly that value.

Table 10.1: Implied ranges of arguments.

argument	input	meaning
gross.value	x	$[x * \text{allowance}, x]$
net.value	see Section 4.2	
long.value	x	$[x * \text{allowance}, x]$
short.value	x	$[x * \text{allowance}, x]$
turnover	x	$[0, x]$
ntrade	n	$[0, n]$
port.size	n	$[1, n]$
alpha.constraint	x	$[x, \text{Inf})$
var.constraint	x	$[0, x]$
bench.constraint	x	$[0, x]$
limit.cost	(must give the range)	
close.number	n	$[n, n]$

Table 10.2: The meaning of `threshold` inputs.

threshold constraint	1 column	2 columns	3 columns	4 columns
trade sell	$-x_1$	x_1	x_1	x_1
trade buy	x_1	x_2	x_2	x_2
portfolio short	0	0	NA	x_3
portfolio long	0	0	x_3	x_4

Threshold Inputs

Table 10.2 shows the meaning of the `threshold` argument when given a matrix with each of the allowable number of columns—a plain vector is the same as a one-column matrix. The subscript of the values in the table indicate the column of the value. So x_1 is from the first column and x_2 is from the second column, and so on. A three-column matrix can only be given in the case of a long-only portfolio.

Positions Inputs

Table 10.3 indicates the allowable number of columns for objects given as the `positions` argument, and the meaning of those columns. See Section 4.7 for more details.

10.4 Troubleshooting

This section lists a number of problems and possible solutions.

Utility Problems

The utility stays the same when I change the risk aversion.

There are at least two possibilities:

Table 10.3: The column order for the `positions` argument.

meaning	2 columns	4 columns	8 columns
min portfolio value	x	x	x
max portfolio value	x	x	x
min trade value		x	x
max trade value		x	x
trade sell threshold value			x
trade buy threshold value			x
portfolio short threshold value			x
portfolio long threshold value			x

1. The objective is neither mean-variance nor mean-volatility—it could be maximizing the information ratio, minimizing variance or maximizing return. Check the `objective.utility` component of the output.
2. You are passing in a value for the `utable` argument. In this case you need to set `force.risk.aver` to `TRUE` in order for the `risk.aversion` argument to override the value that is already in the utility table.

I'm minimizing the variance, and the utility is exactly zero.

The problem is almost certainly that the objective is to maximize the information ratio. You can explicitly set the objective.

This does not occur if neither the `expected.return` nor `utable` arguments are given. It can happen if you give a vector of expected returns that are all zero.

Portfolio Problems

The optimizer is suggesting trades in which some of the assets trade a trivial amount.

Use the `threshold` argument— see Section 4.5.

The optimal portfolio has positions that are very small.

Use the `threshold` argument— see Section 4.5.

I'm building a portfolio and the optimizer is having problems getting the monetary value of the portfolio right.

One possibility is that the `max.weight` argument is too small. If there are only a few assets allowed in the portfolio, the sum of allowable maximum weights can be less than 1. The error of `max.weight` being too small will be caught in simple cases. However, it is possible for the test to be fooled by a combination of restrictions on trading.

Another possibility is that the ranges of the money values (`gross.value`, etc.) are too narrow.

The optimizer is not performing a trade that I am insisting upon via the `lower.trade` or `upper.trade` argument to `trade.optimizer`.

Use the `forced.trade` argument—see Section 4.6.

What if I don't have prices?

You can do an asset allocation and effectively get weights back. See page 99.

I'm trying to generate random portfolios with constraints I know work, but it fails.

You can set the S language seed to some number and try the call to `random.portfolio` again. For example:

```
> set.seed(123)
```

10.5 S Language Problems and Solutions

This section provides a little help with some of the aspects of the S language that you might be most likely to bang your head against.

“Some hints for the R beginner” on the Tutorials page of the Burns Statistics website might also be of help.

Creating Matrices

There are a number of places in this document where we create a matrix. This process generally involves either `cbind` or `rbind`.

`cbind` as in “column bind” and `rbind` as in “row bind”. That sounds simple enough—and it is as long as you pay attention to the details:

```
> cbind(A=c(1, 2))
      A
[1,] 1
[2,] 2
> cbind(A=1, 2)
      A
[1,] 1 2
```

Both of these forms can be useful, but they give matrices that are oriented in different directions.

Debugging

[Burns, 2009] contains a slightly expanded introduction to debugging, and there are a number of additional sources. Here is something to get you started.

Suppose we do something and get an error:

```
> update(soft.op1, expected.return=new.alpha)
Error in trade.optimizer(prices = priceten, variance = varten:
  Object "new.alpha" not found
```

Now in this instance it is clear that the problem is the `new.alpha` that we tried to use. Perhaps we mistyped the name, or we haven't created it yet. But the error could be very much more obscure—it could be unable to find an object that we've never heard of, and possibly in a function that we didn't know we were using.

You can always do the following to get a sense of where the problem happened:

```
> traceback()
5: trade.optimizer(prices = priceten, variance = varten,
  ...
4: eval(expr, envir, enclos)
3: eval(call, parent.frame())
2: update.default(soft.op1, expected.return = new.alpha)
1: update(soft.op1, expected.return = new.alpha)
```

Sometimes just seeing the traceback can be enough to understand what went wrong. Other times, no. For those other cases we can actually go into any of those function calls and see the state at the time of the error. But we can only do that if the ground was prepared before the error happened. Generally it is easy enough to reproduce the error.

We want to set the `error` option to save the contents of the function calls and not just what the calls were:

```
> options(error=dump.frames)
> update(soft.op1, expected.return=new.alpha)
Error in trade.optimizer(prices = priceten, variance = varten:
  Object "new.alpha" not found
> debugger()
Message: Error in trade.optimizer(prices = priceten,    :
  Object "new.alpha" not found
Available environments had calls:
1: update(soft.op1, expected.return = new.alpha)
2: update.default(soft.op1, expected.return = new.alpha)
3: eval(call, parent.frame())
4: eval(expr, envir, enclos)
5: trade.optimizer(prices = priceten, variance = varten
Enter an environment number, or 0 to exit Selection: 5
Browsing in the environment with call:
  trade.optimizer(prices = priceten, variance = varten
Called from: debugger.look(ind)
Browse[1]>
```

Here we have selected frame 5 (the call to `trade.optimizer`) and at this point the objects that that function has created are available to us to explore. You can do:

```
Browse[1]> ls()
```

to see what objects are there. When you want to exit, just type `c` (as in continue):

```
Browse[1]> c
```

Chapter 11

Special Instructions

This chapter contains instructions for how to overcome specific problems. For example, how to do something that the program perceives as an error. In general, these are longer versions of error or warning messages that you might receive.

11.1 Special Instruction 1: Long-only when shorts exist

It is an error to declare the portfolio to be long-only when there are short positions in the existing portfolio. If you have short positions currently and you want to create a long-only portfolio, then you should do the following:

- Declare `long.only=FALSE`
- Use `positions` to set the long-only constraint for the portfolio.

For the last item you could alternatively use `forced.trade` and `lower.trade`, but `positions` is the easier approach.

The `positions` argument could be built like:

```
> posarg <- cbind(rep(0, length(prices)), Inf)
> dimnames(posarg) <- list(names(prices), NULL)
```

Then used like:

```
positions = posarg
```

11.2 Special Instruction 2: Benchmark in long-short optimization

Naively using a benchmark with a long-short portfolio is probably not what you want to do. Creating a benchmark in Portfolio Probe (and elsewhere) is equivalent to being short the benchmark in the amount of the gross value.

If you have a portfolio in the genre of 120/20, then a benchmark is conceptually fine but operationally wrong. If you use the benchmark argument, you

will be short 140% of the benchmark with a 120/20 portfolio while you want to be short 100%.

Whether you have a benchmark in the traditional sense or in the sense that you have a set of liabilities, there are at least two choices:

- Put the portfolio of liabilities as a short asset in the portfolio and stop it from being traded.
- Use a variance matrix that is relative to the liabilities.

There are two functions in the `BurStFin` package that may be of use in this situation. The `var.add.benchmark` function allows you to add a portfolio of assets to a variance matrix. `var.relative.benchmark` transforms a variance matrix to be relative to one of its assets.

Chapter 12

Adjusting Optimization Speed and Quality

The default control settings for the optimizer are a compromise between the time it takes to finish and the quality of the solution.

Optimization is usually done in one of two settings:

- Backtesting strategies
- Production

In backtesting, the quality of the solution is much less of an issue than the speed of the optimization. If the default setting seems too slow to you, you can decrease the maximum number of iterations allowed. Setting

```
iterations.max = 5
```

is a possibility—this is not unreasonably small for a lot of problems.

For production where time is not critical and quality is more of concern, then an easy way to get better solutions is to add the argument:

```
stringency = 1
```

This will take a few times longer to execute, but will give at least as good of a solution.

The rest of the chapter goes into more detail about the options that you have and what they mean.

12.1 Staying at a Given Solution

A special type of “optimization” is to just return an object where the solution is the same as the starting value. At first glance this seems like a useless operation, but in fact there are many reasons to do this—here are a few:

- See the cost of a solution

- See constraint violations of a solution
- Get expected returns, variances and so on of random portfolios (this is what `randport.eval` does)
- Ensure that another program is doing the same problem
- Examine the utility of your current portfolio, or of a proposed trade

`trade.optimizer` makes this easy. Give the trade that you want to test as the `start.sol` argument, and set `funeval.max` to zero or one. An asset allocation example is:

```
> aa.gs1 <- trade.optimizer(aprice, avar.medium,
+   aret.medium, gross.value=ten.thou,
+   long.only=TRUE, utility="mean-variance",
+   risk.aversion=.05, start.sol=e60b40, funeval=0)
> valuation(aa.gs1)$weight
      equities      bonds
      0.6        0.4
> aa.gs1$utility.value
[1] 2.08
> aa.gs1$alpha.value
[1] 6.4
> aa.gs1$var.value
[1] 169.6
```

Note that setting:

```
iterations.max = 0
```

is not the same. This still allows the pre-iteration and post-iteration operations and the result could be substantially different from the starting solution.

If the existing portfolio is where you want to stay, then you can use the argument:

```
ntrade = 0
```

12.2 Reducing Time Use

You may want to reduce the time that an optimization takes because you are doing a large number of optimizations. The only viable way to reduce time is to reduce the `iterations.max` argument.

For a lot of problems 5 iterations is likely to optimize enough that backtests will give the right indication. Zero iterations will almost surely be too few, but even one iteration may not be so terrible as to be useless.

12.3 The Optimization Process

This section briefly describes the optimization algorithm so that the suggestions for improving quality in the next section will make sense.

The optimization process can be broken into three levels:

- sub-algorithms
- iterations
- runs

The algorithm for `trade.optimizer` is generally described as a genetic algorithm. That is a simplification. Each iteration performs several sub-algorithms—some of them genetic (plus simulated annealing), and others can be described as greedy algorithms.

A run is a series of iterations. The run ends when too many consecutive iterations fail to find an improvement. That number of failures is controlled by the `fail.iter` argument, which has a default of zero. So, by default, any iteration that fails to improve will cause the run to end. If `fail.iter` is 1, then it takes two consecutive iteration failures to end the run.

The default is to perform a single run. Multiple runs are done if the `stringency` argument is positive. The stringency criterion is satisfied if the best trade has been found again (in different runs) `stringency` times.

For example, suppose `stringency` is 1. After each run we have the best found so far. If the result of the new run is the same (the same assets traded, the same number of units traded) as the best found previously, then the stringency criterion is met and the process stops. If `stringency` were 2, then it would take three runs with the same result to end the process.

When stringency is positive, there are three types of runs:

- initial runs
- final runs
- non-convergence run

Initial runs perform the optimization as stated. The number of initial runs is determined by the `runs.init` control argument.

If the stringency hasn't been satisfied in the initial phase, then the final runs are started. The key difference is that only assets that were involved in trades in the initial phase are eligible to trade in this phase. The number of final runs is controlled by `runs.final`.

If the stringency is still not satisfied at the end of the final runs, then one more run is done. This run starts with the best solution found in any of the previous runs. The restriction to trading as in the final runs remains in place. The maximum number of iterations allowed in this run is a multiple of the maximum number of iterations for the other runs—that multiple is given by the `nonconverge.mult` control argument.

12.4 Improving Quality

The easiest—and generally most effective—approach to improving quality is to set `stringency` to 1.

You can improve quality by increasing `iterations.max` and `fail.iter`.

If the `converged` component of the result is `FALSE`, then it is probably the case that the best solution has not been found. You can restart the optimization from where it left off:

```
> opt1 <- trade.optimizer(prices, varian, ...)
> opt2 <- trade.optimizer(prices, varian,
+   start.sol=opt1, ...)
```

For more thorough optimization, increase `iterations.max` and `fail.iter`:

```
> opt3 <- trade.optimizer(prices, varian, start.sol=opt1,
+   iterations=100, fail=10, ...)
```

Another way of doing this same thing would be:

```
> opt3 <- update(opt1, start.sol=opt1, iterations=100,
+   fail=10)
```

S code note

If an object is a list and has a component named `call` (that is an image of the command that created the object), then the `update` function will recompute the `call` and change any arguments that are given in the `call` to `update`. The `update` function can be used on the results of `trade.optimizer` and `random.portfolio`.

A natural thing to do if you want a quality solution is to increase the maximum number of iterations to something like 100 or 1000. That is not the best approach. For very large problems 100 iterations might be useful, but increasing the number of runs and increasing the number of iterations slightly is usually a better approach.

12.5 Testing Optimization Quality

The optimization algorithm is random (pseudo-random to be technical). You get a consistent answer for any particular problem because the same random seed is used by default. You can get different paths (and hence probably different answers) by setting the seed argument to `NULL`. One way of getting a new optimization is:

```
new.optobj <- update(orig.optobj, seed=NULL)
```

You can create a list of a number of solutions for comparison with commands similar to:


```
> opt.list <- vector("list", 10)
> for(i in 1:10) opt.list[[i]] <- update(orig.opt, seed=NULL)
> summary(unlist(lapply(opt.list, function(x)
+   x$results["objective"]))), digits=7)
> summary(unlist(lapply(opt.list, function(x)
+   trade.distance(x, orig.opt))), digits=7)
```

Here we've created an empty list, and filled it up with solutions to the problem. Finally we looked at the distribution of the achieved objective, and the distribution of trade distances from the original solution.

S code note

Some care needs to be used when doing `seed=NULL` in a call to `update`. If `seed` was used in the original call (to `trade.optimizer`), then the result will be to use the default value of `seed` (which is fixed) rather than creating a new random seed.

The `update` function creates its own call to the function and then evaluates that call. Saying `arg=NULL` removes `arg` from the call if it was in the original call.

Chapter 13

Utility

Utility is ignored in random portfolios (but see Section 3.7), however it is central to optimization.

The descriptions of the utilities use a few entities:

- The weight vector w .
- The vector of expected returns α .
- The variance matrix V .
- The transaction costs $C(w)$.

In long-only portfolios all of the elements of w are non-negative and they sum to 1. In long-short portfolios the weights may be negative, and it is the absolute values of the weights that sum to 1—that is, the weight of an asset is its value in the portfolio divided by the gross value of the portfolio.

$C(w)$ is actually an abuse of notation for the transaction cost. Transaction cost obviously depends on the existing portfolio. Cost is not strictly a function of weight unless the gross value for the portfolio is constant.

13.1 Maximum Information Ratio

This utility (which is the default if there is both a variance and expected returns) is specified with the argument:

```
utility = "information ratio"
```

or

```
utility = "exocost information ratio"
```

The generic information ratio problem is to maximize:

$$\frac{\alpha^T w}{\sqrt{w^T V w}}$$

over all vectors w that satisfy the constraints.

This inherently assumes zero costs. There are two approaches (at least) to incorporating costs. The first is to subtract the trading cost from the expected value:

$$\frac{\alpha^T w - C(w)}{\sqrt{w^T V w}}$$

This is what you get with:

```
utility = "information ratio"
```

The second approach is to have a separate term for costs:

$$\frac{\alpha^T w}{\sqrt{w^T V w}} - C(w)$$

You get this second version with:

```
utility = "exocost information ratio"
```

13.2 Mean-Variance Utility

Mean-variance utility is invoked with:

```
utility = "mean-variance"
```

With mean-variance utility we maximize:

$$\alpha^T w - \gamma w^T V w - C(w)$$

where γ is the risk aversion parameter and w satisfies all of the constraints.

Many implementations have a one-half in the variance term, meaning that the risk aversion parameter in those cases is a factor of two different than here. Also in some optimizers the parameter used divides the variance rather than multiplies it—in which case it is a risk tolerance parameter.

The mean-variance risk aversion parameter is invariant to annualization. As long as both `expected.return` and `variance` are for the same time scale, it doesn't matter what time scale it is. However risk aversion is affected when `expected.return` and `variance` are for returns that are put into percent. When percent returns are used, the risk aversion is divided by 100.

[Kallberg and Ziemba, 1983] classify risk aversion greater than 3 as very risk averse, 1 to 2 as moderate risk aversion, and less than 1 as risky—these numbers would be divided by 100 for data representing percent returns. Some additional sense of suitable values for the risk aversion parameter may be found in [Burns, 2003b].

All of the above are focused on long-only portfolios. In long-short portfolios the variance is a smaller quantity so it takes a larger risk aversion to have the same effect. In general risk aversion is likely to be larger in long-short situations than for long-only.

13.3 Mean-Volatility Utility

Mean-volatility utility is invoked with:

```
utility = "mean-volatility"
```

With mean-volatility utility we maximize:

$$\alpha^T w - \zeta \sqrt{w^T V w} - C(w)$$

where ζ is the risk aversion parameter and w satisfies all of the constraints.

The mean-volatility risk aversion does depend on annualization. For example the risk aversion for annual data is $\sqrt{252}$ times the risk aversion for daily data. But volatility aversion is invariant to whether or not the data are in percent.

13.4 Minimum Variance

Minimum variance utility is specified with:

```
utility = "minimum variance"
```

The minimum variance utility minimizes:

$$w^T V w + C(w)$$

over the w that satisfy the constraints.

This utility is used when:

- The `utility` argument is set to be "minimum variance".
- The `expected.return` argument is not given (or is NULL) and distance is not being minimized.
- The utility is either mean-variance or mean-volatility and `risk.aversion` is `Inf`.

13.5 Maximum Expected Return

The maximum expected return utility is specified via:

```
utility = "maximum return"
```

With this utility we maximize:

$$\alpha^T w - C(w)$$

where w must satisfy the constraints.

This utility is used when:

- The `utility` argument is set to be "maximum return".

- The **variance** argument is not given (or is NULL) and distance is not being minimized.
- The utility is mean-variance or mean-volatility and **risk.aversion** is 0.

In the last of these items, the stated utility in the output remains mean-variance or mean-volatility, but a look at the definitions shows that it is really maximum expected return.

13.6 Minimum Distance

This is specified with:

```
utility = "distance"
```

This minimizes the distance from the portfolio to the specified target portfolio plus the trading costs. More details are found in Section 6.4.

13.7 Going Farther

- Chapter 14 discusses utilities where multiple variances and/or expected return vectors are given.

Chapter 14

Advanced Features

This chapter will be of use if you have or want:

- Multiple variances
- Multiple expected return vectors
- Multiple benchmarks
- Compact variances (seldom recommended)

If you merely want constraints on multiple benchmarks, that is easy and is discussed on page 72.

14.1 Multiplicity

The computation that is performed is really controlled by three entities: the alpha table, the variance table and the utility table. In the standard case of no multiplicity, these entities can safely be ignored. When you do have multiplicity, then understanding what these objects are saying is a good idea. If you understand them, then you can check to make sure that it is doing what you want done. If it isn't doing what you want automatically, then you can tell it to do what you do want.

If you are generating random portfolios but not optimizing, then you can skip parts of this chapter. The recommended route in this case is:

- Section 14.2 Alpha and Variance Tables
- Section 14.3 Variance Constraints
- Section 14.4 Expected Return Constraints

Table 14.1 describes arguments to `trade.optimizer` (and `random.portfolio`) that are useful when dealing with multiplicity. All of these arguments have defaults—you only need to give values for these if the default behavior is not what you want.

Table 14.1: Arguments for multiple variances and expected returns.

Argument	Description
<code>vtable</code>	gives the variance-benchmark combinations
<code>atable</code>	gives the expected return-benchmark combinations
<code>utable</code>	describes all the combinations to get utilities
<code>quantile</code>	states which quantile of the multiple objectives to use
<code>dest.wt</code>	gives weights to the multiple objectives if desired

14.2 Alpha and Variance Tables

The purpose of the alpha table is to state what combinations of expected return vectors and benchmarks will be used in the problem. Likewise the variance table identifies variance-benchmark combinations that will be used.

An alpha table is a matrix with two rows and an arbitrary number of columns. In simple cases the alpha table will look like:

```
> optobj$atable
      [,1]
alpha    0
benchmark -1
attr(,"benchmarks")
[1] ""
```

Or if a benchmark is given, it will look something like:

```
> optobj$atable
      [,1]
alpha    0
benchmark 500
attr(,"benchmarks")
[1] "spx"
```

The first row gives the zero-based index of the expected return. If there is only one set of expected returns, then this is always zero. If there are no expected returns, this is negative one.

The second row gives the zero-based index of the corresponding benchmark. A negative number in this row means there is no benchmark.

There is a `benchmarks` attribute which is a character representation of the benchmarks. When you are passing an object in as the `atable` argument, the second row of the matrix itself is overwritten with the information in the `benchmarks` attribute (which must be present).

Variance tables are essentially the same as alpha tables. The first row, of course, corresponds to variances rather than alphas. The other change is that there is a third row which states if this combination is only used in the utility (value 1) or if it is used in constraints as well (value 0). This information is used to more efficiently generate random portfolios. As with alpha tables when you pass an object in as the `vtable` argument, the second row is overwritten with the information given in the `benchmarks` attribute of the object.

Here is an example of the alpha table and variance table from an optimization that has two variance matrices, three expected return vectors, and no benchmarks:

```
> opt1$atable
      [,1] [,2] [,3]
alpha    0   1   2
benchmark -1  -1  -1
attr("benchmarks")
[1] "" "" ""
> opt1$vtable
      [,1] [,2]
variance    0   1
benchmark  -1  -1
utility.only 1   1
attr("benchmarks")
[1] "" ""
```

Here are the tables from an optimization with the same variance and expected returns but with two benchmarks:

```
> opt2$atable
      [,1] [,2] [,3] [,4] [,5] [,6]
alpha    0   1   2   0   1   2
benchmark  8   8   8   9   9   9
attr("benchmarks")
[1] "Index1" "Index1" "Index1" "spx" "spx" "spx"
> opt2$vtable
      [,1] [,2] [,3] [,4]
variance    0   1   0   1
benchmark  8   8   9   9
utility.only 1   1   1   1
attr("benchmarks")
[1] "Index1" "Index1" "spx" "spx"
```

note

The values in the alpha table, variance table and utility table are zero-based because these entities are sucked directly into C where indexing is zero-based rather than one-based as in the S language.

14.3 Variance Constraints

This section discusses the argument:

- `var.constraint`

Simple use of this argument is discussed on page 145. This argument also handles problems that are too complex for the `bench.constraint` argument.

The `var.constraint` argument can be:

- a plain vector (with one or more values) with or without names
- a one-column matrix with or without row names
- a two-column matrix with or without row names

Plain vectors are equivalent to one-column matrices. One-column matrices are equivalent to two-column matrices where the first column is all zeros.

The names on vectors or row names of matrices specify the zero-based index of the column of the variance table. When there are no names, then the first column(s) are implied (in order).

So all six of the following are specifying the same constraints:

```
var.constraint = vc.vec
var.constraint = vc.1mat
var.constraint = vc.2mat
var.constraint = vc.vec.nam
var.constraint = vc.1mat.nam
var.constraint = vc.2mat.nam
```

where:

```
> vc.vec
[1] 2.4 1.5
> vc.1mat
  [,1]
[1,] 2.4
[2,] 1.5
> vc.2mat
  [,1] [,2]
[1,]  0 2.4
[2,]  0 1.5
> vc.vec.nam
  0 1
2.4 1.5
> vc.1mat.nam
  [,1]
0 2.4
1 1.5
> vc.2mat.nam
  [,1] [,2]
0  0 2.4
1  0 1.5
```

All of these are saying the variance-benchmark combination in the first column of the variance table has an upper bound of 2.4 and the variance-benchmark combination in the second vtable column is to be no more than 1.5.

If you were to give the equivalent of `vc.vec.nam` directly, you would do:

```
var.constraint = c('0'=2.4, '1'=1.5)
```

A very similar command means something quite different:

```
var.constraint = cbind(2.4, 1.5)
```

means that the (first) variance is constrained to be at most 1.5 but no less than 2.4. (In this case you will get an error, but you can not depend on getting an error when you confuse the direction of matrices.)

Two-column matrices give lower bounds as well as upper bounds on the variance. Lower bounds are most likely to be useful in generating random portfolios, but could potentially be of interest in optimization as well.

Consider the matrix:

```
> varcon.m1 <- rbind(c(1.2, 1.4), c(1.3, 1.5))
> varcon.m1
  [,1] [,2]
[1,] 1.2 1.4
[2,] 1.3 1.5
```

Suppose that there are three columns in `vtable` (perhaps the `variance` argument is a three-dimensional array with 3 slices in the third dimension). Then the command:

```
var.constraint = varcon.m1
```

would say that the first variance-benchmark combination is restricted to be between 1.2 and 1.4, the second variance-benchmark is restricted to be between 1.3 and 1.5, and the third is unrestricted.

Now let's add some row names.

```
> varcon.m2 <- varcon.m1
> dimnames(varcon.m2) <- list(c(2, 0), NULL)
> varcon.m2
  [,1] [,2]
2  1.2  1.4
0  1.3  1.5
```

The command:

```
var.constraint = varcon.m2
```

says that the third variance-benchmark is to be between 1.2 and 1.4, the first variance-benchmark is restricted to be between 1.3 and 1.5, and the second variance-benchmark is unconstrained.

Some examples with multiple variances can be found on page 153.

14.4 Expected Return Constraints

This section discusses the argument:

- `alpha.constraint`

Simple use of this argument is discussed on page 69.

Use of this argument is analogous to the advanced use of `var.constraint` (page 145) except that a one-column matrix is a lower bound rather than an upper bound, and indices refer to expected return-benchmark combinations (columns of `atable`) rather than to variance-benchmark combinations.

So for instance:

```
alpha.constraint = c("1"=5.6, "0"=2.3)
```

is equivalent to

```
alpha.constraint = ac.2mat.nam
```

or

```
alpha.constraint = ac.2mat
```

where:

```
> ac.2mat.nam
  [,1] [,2]
1  5.6  Inf
0  2.3  Inf
> ac.2mat
  [,1] [,2]
[1,]  2.3  Inf
[2,]  5.6  Inf
```

14.5 Multiple Utilities

This section does not pertain to generating random portfolios (unless a utility constraint is imposed).

When there are multiple columns in the alpha or variance tables, there are usually multiple utilities for any given trade. In order to do an optimization, we need to combine that set of numbers into just one number.

A common approach is to get a min-max solution. This process consists of always picking out the worst answer for each trade—we then try to get the best of these worst values.

Another approach would be to take the median of the values for a trade. This represents an optimistic attitude, while min-max is very pessimistic—min-max assumes that nature will do its worst against us. The `quantile` argument allows us to choose either of these, or something in between. The useful range of values for `quantile` is 0.5 (the median) to 1 (the maximum—yielding the min-max solution).

Each of the individual utilities for a trade is called a “destination”. What goes into a destination is under user control. For example, if some destinations are deemed to be more important than others, they can be given more weight with the `dest.wt` argument.

There are two sorts of destination weight. Each utility has a weight within its destination. The within destination weights can be used to create a weighted average of utilities within a single destination; they can also be used to modify the utility as demonstrated in the dual benchmark example which starts on page 150.

The weights in the `dest.wt` argument, on the other hand, are weights among the destinations which are used along with the quantile argument to control the selection of the objective from among the destinations. For example in scenario analysis (see page 156) `dest.wt` could be the probability assigned to each scenario.

note

The answer when `dest.wt` is not given is, in general, slightly different than when `dest.wt` is given and all of the weights are equal.

The actual problem that is done is mostly specified by the utility table.

14.6 Utility Tables

The utility table is a matrix that gives the combinations of variances, expected returns and benchmarks that are to be used and what to do with them. Here is the utility table from our example of two variances and three expected returns:

```
> opt1$utable
      [,1] [,2] [,3] [,4] [,5] [,6]
alpha.spot      0   1   2   0   1   2
variance.spot   0   0   0   1   1   1
destination     0   1   2   3   4   5
opt.utility     1   1   1   1   1   1
risk.aversion   1   1   1   1   1   1
wt.in.destination 1   1   1   1   1   1
```

The first row is the expected return—this is the zero-based index of the columns of `atable`. That is, it indicates a particular expected return-benchmark combination. (But note that this is the zero-based location of the distance for columns corresponding to distance utilities.)

The second row is similar to the first except that it is columns of the `vtable`.

The third row is the zero-based index of the destination for the utility with this combination of alpha and variance. The utility for the combination is computed, multiplied by its weight in the destination (the sixth row of the utility table) and added to whatever is there. That is, a destination may hold a weighted average of utilities. If a destination is a negative number, the utility is not placed into a destination—these are combinations that are presumably to be used in a constraint. There must be at least one zero in the destination row, and the destinations must be numbered from zero upwards with no integers skipped.

The fourth row is an indicator for the type of utility to be performed. The allowable codes are given in Table 14.2.

Table 14.2: Utilities codes for the utility table.

Code	Meaning
0	mean-variance
1	information ratio
2	exocost information ratio
3	minimum variance
4	maximum return
5	mean-volatility
6	distance

The fifth row holds the risk aversion parameter for the cases where it is used—the value is ignored in the other cases. When the `utable` argument is given, then the risk aversion in the utility table is used rather than using the value of the `risk.aversion` argument unless `force.risk.aver` is set to `TRUE`.

The example utility table above specifies that six utilities are to be computed and each put into a different destination. Each combination of expected return and variance is used, and the information ratio is the utility in all cases.

14.7 Multiplicity Examples

This section provides a few examples of using multiple variances, expected returns and benchmarks.

Dual Benchmarks

Dual benchmarks can be used, for example, to incorporate predictions about a future rebalancing of the benchmark. Due to the relative movements that have occurred since the last rebalance, we may have predictions of what assets will enter and exit the benchmark and with what weight. The portfolio can be optimized against both benchmarks—this allows the portfolio to be rebalanced to some extent against the new benchmark while still having protection relative to the current one. This can mean a cheaper rebalance if it is done before others do their rebalancing.

If our benchmark is `spx` and our prediction of the post-rebalance benchmark is `newspx`, then minimizing variance relative to the two benchmarks can be done as:

```
> opts1 <- trade.optimizer(prices, varian,
+   long.only=TRUE, gross.value=1e6, existing=cur.port,
+   utility = "minimum variance",
+   benchmark=c("spx", "newspx"), quantile=1)
```

Often a min-max solution is found for dual benchmarks. To get such a solution, set `quantile` to 1. The min-max solution says that we want to minimize the worst tracking error. It is typical (though not always true) that the utility in the optimal solution is the same for both benchmarks:

```

> sqrt(252 * opts1$var.values)
[1] 0.009029673 0.009029442
> opts1$utility.values
[1] 3.235515e-07 3.235350e-07

```

In this example the tracking errors are the same at 90 basis points, and the utilities are equal as well.

A reasonable complaint about this optimization is that it treats both benchmarks equally even though the new benchmark is speculative while the current one is known precisely. It is probably more reasonable to insist on a smaller tracking error against the current benchmark.

A difference in the tracking errors is easily achieved with the addition of another argument. The utility table needs to be changed slightly from its default. So we recover the utility table from the original optimization, make the change, then do a new optimization:

```

> utab1 <- opts1$utable
> utab1
           [,1] [,2]
alpha.spot      0    0
variance.spot   0    1
destination     0    1
opt.utility     0    0
risk.aversion   1    1
wt.in.destination 1    1
> utab1[6,1] <- 1.4
> utab1
           [,1] [,2]
alpha.spot    0.0  0
variance.spot 0.0  1
destination   0.0  1
opt.utility   0.0  0
risk.aversion 1.0  1
wt.in.destination 1.4 1

```

All that we have done is change the weight-in-destination value for the first (current) benchmark from 1 to 1.4. Now the optimization yields:

```

> opts2 <- trade.optimizer(prices, varian,
+   long.only=TRUE, gross.value=1e6, existing=cur.port,
+   utility = "minimum variance",
+   benchmark=c("spx", "newspx"), quantile=1,
+   utable=utab1)
> sqrt(252 * opts2$var.val)
[1] 0.008227913 0.009735447
> opts2$utility.val
[1] 3.761031e-07 3.761069e-07

```

The utility values are still equal, but now the utility represents something different for the current benchmark than for the new benchmark, so the tracking

errors are different—now they are 82 basis points for the current benchmark and 97 for the new benchmark.

There is not a limit on the number of benchmarks that you can use—just add the names to the vector given as the `benchmark` argument.

If you have an active portfolio, you may want to put on constraints relative to each benchmark. This is done precisely the same as with a single benchmark constraint. Adding a second benchmark constraint is just like adding any other constraint—it does not create multiple utilities. An example is given on page 95.

Benchmark-relative Utility and Absolute Variance Constraint

We saw (on page 94) that it is easy to have a utility in absolute terms on the portfolio with a benchmark constraint. The reverse problem—benchmark-relative utility with a constraint on portfolio volatility—is just as easy to state, but not as easy to do.

We need to create a variance table to pass in as the `vtable` argument. It needs a column with the benchmark to use in the utility and a second column without the benchmark for the constraint. We can do this with:

```
> vtab.buav <- rbind(c(0,0), c(0,-1), c(1,0))
> attr(vtab.buav, "benchmarks") <- c("spx", "")
```

We will also need to create a utility table, but in this case it is simple enough to create on the fly within the call to the optimizer. Our optimizer call can look something like:

```
> opt.buav <- trade.optimizer(prices, varian,
+   expected.return=alphas, benchmark='spx',
+   var.constraint=c('1'=100), vtable=vtab.buav,
+   utable=rep(c(0,1), c(3,3)), ... )
```

We can check the tables that are created to see if it is doing what we want.

```
> opt.buav$aatable
      [,1]
alpha      0
benchmark 500
attr(,"benchmarks")
[1] "spx"
> opt.buav$vtable
      [,1] [,2]
variance      0  0
benchmark    500 -1
utility.only   1  0
attr(,"benchmarks")
[1] "spx" ""
> opt.buav$utable
      [,1]
alpha.spot      0
variance.spot   0
```



```

destination      0
opt.utility      1
risk.aversion    1
wt.in.destination 1

```

The utility table we created says to maximize the information ratio. We need to put a name on the object given as the variance constraint because we want to constrain the variance-benchmark combination given in the second column of the variance table (not the first column).

Rival Variance Forecasts

If you have two or more variance matrices, you can use these in a single optimization. For example if you have a statistical factor model, a fundamental factor model and a macroeconomic factor model, you could use all three in the optimization. Since these are created using different sources of information, it is reasonable to suppose that results may be better by using more than one of them.

The first task is to create a suitable variance object, which will (almost always) be a three-dimensional array:

```

> varmult <- array(c(vstat[vnam, vnam], vfund[vnam, vnam],
+   vmacro[vnam, vnam]), c(length(vnam), length(vnam), 3),
+   list(vnam, vnam, NULL))

```

S code note

A three-dimensional array is a generalization of a matrix—instead of a dim attribute that has length 2, it has a length 3 dim. Likewise, the dimnames is a list that has three components.

We need each of the variance matrices to be in a slice of the third dimension. In the example above, we are ensuring that the matrices we are pasting together all correspond to the same assets in the same order by subscripting with `vnam`.

You need to decide which quantile to optimize. The best quantile to use probably depends on how the variance matrices relate to each other and on the type of optimization that you are performing. Once this decision is made, the optimization is performed as with a single variance:

```

> opts.m1 <- trade.optimizer(prices, varmult,
+   long.only=TRUE, gross.value=1e6, expected.return=alphas,
+   bench.constrain = c(spx=.05^2/252), quantile=.7)

```

This command is maximizing the information ratio with a tracking error constraint as is done on page 94. In this case the tracking error is constrained relative to each of the variance matrices. The information ratio that is being optimized is the 70% quantile of the set of information ratios using each of the variances.

If you are doing a variance constraint, for instance maximizing the information ratio in a long-short portfolio as on page 103, then the approach is slightly different. You need to make the variance constraints argument the proper length yourself:

```
> opts.m2 <- trade.optimizer(prices, varian,
+   expected.return=alphas, gross.value=1e6,
+   net.val=c(-1e4, 2e4), ntrade=55, max.weight=.05,
+   var.constraint=rep(.04^2/252, 3))
```

If the value given as the variance constraint weren't replicated, then only the first variance would be constrained.

The `var.constraint` argument is the more general. Benchmark constraints actually create variance constraints—the `bench.constraint` argument only exists for convenience. (When more than one variance is given, `bench.constraint` puts the constraint on each of the variances.) More control can be gained over variance constraints by putting names on the vector given. The names need to be the zero-based index of the columns of the variance table (see page 145). For example to constrain the first variance to a volatility of 4% and the third to a volatility of 5%, the command would be:

```
> opts.m3 <- trade.optimizer(prices, varian,
+   expected.ret=alphas, gross.val=1e6,
+   net.val=c(-1e4, 2e4), ntrade=55, max.weight=.05,
+   var.constraint=c("0"=.04^2/252, "2"=.05^2/252))
```

S code note

The names given inside the `c` function need to be put in quotes because they are not proper names of S objects. S would try to interpret them as numbers if they were not in quotes.

If any element of the `var.constraint` argument is named, then they all need to be. When there are no names, the columns of the variance table are used in order.

A standard use of multiple variances is to optimize with one, and check the tracking error or volatility with the other. One good reason to evaluate the solution on a variance not used in the optimization is because of the bias created in the optimization process—the tracking error or volatility that is achieved in the optimization is a downwardly biased estimate of what will be realized. Using multiple variances in the optimization will reduce the amount of bias in any particular variance matrix, though some bias will still remain.

Multiple Time Periods

You may have forecasts for multiple time periods. For example, you may have GARCH forecasts of the variance matrix for several different time horizons. You want the portfolio to do well at all of the time scales. A key issue is how to weight the different time scales. You may want different risk profiles at different time scales—for instance, expected returns may be given relatively more weight versus risk for short time horizons.

Credit Risk

The variance matrices that we have talked about have all been measures of market risk. There are other sources of risk—credit risk, for example—that could be included in an optimization. For more on credit risk, see for example [Gupton et al., 1997]. While we’ll speak of credit risk here, keep in mind that other forms of risk could be used instead or in addition.

Suppose that we have `var.market` and `var.credit` that represent the two types of risk. While it may be possible to scale the credit risk so that it is comparable to the market risk, it may be easier to accept that they are different. If the two forms of risk are not comparable, then you can perform the optimization with the market risk as you normally would and put a constraint on the credit risk.

When deciding how to constrain the credit risk, the first thing to do is explore its limits. One end of the range is when credit risk is ignored entirely in the optimization, the other end is when the credit risk is minimized with no concern for the market risk.

```
> op.cred0 <- trade.optimizer(prices, var.market, ...)
> trade.optimizer(prices, var.credit, ...,
+   funeval=0, start.sol=op.cred0)$var.value
[1] 26.02844
> trade.optimizer(prices, var.credit,
+   utility="minimum variance", ...)$var.value
[1] 11.92467
```

So the range of interest is about 12 to 26. If we decided to constrain the credit risk to 18, then we would do:

```
> var.mult <- array(c(var.market[anam, anam],
+   var.credit[anam, anam]), c(length(anam),
+   length(anam), 2), list(anam, anam, NULL))
> op.cred1 <- trade.optimizer(prices, var.mult, ...,
+   var.constraint=c("1"=18))
```

The first command creates the three dimensional array containing the two types of risk (taking care to make sure the data line up properly). The second command does the optimization. The variance constraint is a named vector where the name is the zero-based index to the third dimension of the variance array. In this case the name is “1” because we are constraining the second variance.

But `op.cred1` is not the optimization that we are aiming for. The credit risk is used in the utility as well as being constrained. We need to tell the optimizer not to use credit risk in the utility by passing in a `utable` argument. We can make a simple change to the utility table from the optimization we’ve just done.

```
> op.cred1$utable
      [,1] [,2]
alpha.spot      0  0
variance.spot   0  1
destination     0  1
opt.objective   1  1
```

```

risk.aversion      1    1
wt.in.destination  1    1
> uticred <- op.cred1$utable
> uticred[3,2] <- -1
> uticred
[,1] [,2]
alpha.spot         0    0
variance.spot      0    1
destination        0   -1
opt.objective      1    1
risk.aversion      1    1
wt.in.destination  1    1
>
> op.cred2 <- trade.optimizer(prices, var.mult, ...,
+   var.constraint=c("1"=18), utable=uticred)

```

The change we made was to say that the column with credit risk in it shouldn't go into a destination. One indication that we are doing it right is that the `utility.values` component of `op.cred2` has length one, while it is length two for `op.cred1`. Since we only want one variance in the utility, the utility that we want will have length one.

Multiple Scenarios

Scenario analysis creates some hypotheses about what might happen in the future, and then tries to pick the best course of action. The assumption is that the future will look like one of the scenarios, or perhaps some combination of scenarios.

Our first step is to create some data. We will create a “medium” scenario, a “crash” scenario and a “prosperous” scenario. (The probability that these numbers are reasonable is close to zero.)

```

> anam4 <- c("equities", "bonds", "commodities", "cash")
> acor.medium <- matrix(c(1, .2, .1, .2, 1, .1, .1, .1, 1),
+   3, 3)
> acor.medium
           equities bonds commodities
equities   1.0  0.2      0.1
bonds      0.2  1.0      0.1
commodities 0.1  0.1      1.0
> acor.crash <- matrix(c(1, -.3, .3, -.3, 1, -.2, .3,
+   -.2, 1), 3, 3)
> acor.crash
           [,1] [,2] [,3]
[1,]  1.0 -0.3  0.3
[2,] -0.3  1.0 -0.2
[3,]  0.3 -0.2  1.0
> acor.prosper <- matrix(c(1, .4, .3, .4, 1, .2, .3,
+   .2, 1), 3, 3)

```

```

> acor.prosper
      [,1] [,2] [,3]
[1,]  1.0  0.4  0.3
[2,]  0.4  1.0  0.2
[3,]  0.3  0.2  1.0
> avol.medium <- c(20, 8, 12)
> avol.crash <- c(35, 16, 20)
> avol.prosper <- c(15, 5, 7)
> avar.medium <- rbind(cbind(t(acor.medium * avol.medium) *
+   avol.medium, 0), 0)
> avar.crash <- rbind(cbind(t(acor.crash * avol.crash) *
+   avol.crash, 0), 0)

```

S code note

The last command creating `avar.crash` does several things at once. First (in the center of the command) it multiplies the correlation matrix times the volatility vector which multiplies each row of the correlation matrix by the corresponding volatility. Then it transposes the matrix and multiplies again by the volatility vector—this is now multiplying what originally were the columns of the correlation matrix by their corresponding volatilities. Then it binds on a column of all zeros (for cash), and finally binds on a row of all zeros. Below you can see the result of all this manipulation once the asset names are put onto the variance matrix.

```

> avar.prosper <- rbind(cbind(t(acor.prosper * avol.prosper)
+   * avol.prosper, 0), 0)
> dimnames(avar.medium) <- list(anam4, anam4)
> dimnames(avar.crash) <- list(anam4, anam4)
> dimnames(avar.prosper) <- list(anam4, anam4)
> avar.medium
      equities bonds commodities cash
equities    400  32.0      24.0    0
bonds        32  64.0       9.6    0
commodities   24   9.6     144.0    0
cash          0   0.0       0.0    0
> avar.crash
      equities bonds commodities cash
equities   1225 -168      210    0
bonds     -168  256      -64    0
commodities 210  -64      400    0
cash         0   0       0    0
> avar.prosper
      equities bonds commodities cash
equities   225.0  30      31.5    0
bonds      30.0  25       7.0    0
commodities 31.5   7      49.0    0
cash        0.0   0       0.0    0
> aret.medium <- c(8, 4, 5, 3)

```

```

> aret.crash <- c(-20, 6, -9, 3)
> aret.prosper <- c(25, 3.5, 8, 3)
> names(aret.medium) <- anam4
> names(aret.crash) <- anam4
> names(aret.prosper) <- anam4
> aret.crash
  equities    bonds commodities    cash
      -20         6         -9         3
> aret.prosper
  equities    bonds commodities    cash
     25.0     3.5      8.0      3.0

```

Since we are doing asset allocation, there are no prices for the assets and we want the answer to be the weights of the assets. Hence we want to create a “price” vector that is all ones and set the gross value to a useful value.

```

> aprice <- rep(1, 4)
> names(aprice) <- anam4
> aprice
  equities    bonds commodities    cash
         1         1         1         1
> ten.thou <- 10000 + c(-.5, .5)
> op.medium <- trade.optimizer(aprice, avar.medium,
+   aret.medium, gross.value=ten.thou, long.only=TRUE,
+   utility="mean-variance", risk.aversion=.02)
> op.medium$new.portfolio / 100
  equities    bonds commodities    cash
     27.86     20.83     28.69     22.62

```

Above we have performed a mean-variance optimization with the “medium” data, and then shown the weights of the resulting portfolio in percent. (The variance matrix is not positive definite but the optimizer doesn’t care—note, though, that no more than one zero variance should be in a single variance matrix.)

In order to perform simultaneous optimization, we need to put the three variance matrices into a three-dimensional array, and the expected return vectors into a matrix:

```

> avar.all <- array(c(avar.crash, avar.medium, avar.prosper),
+   c(4, 4, 3))
> dimnames(avar.all) <- list(anam4, anam4, NULL)

```

S code note

A three-dimensional array is a generalization of a matrix—instead of a `dim` attribute that has length 2, it has a length 3 `dim`. Likewise, the `dimnames` is a list that has three components.

We need each of the variance matrices to be in a slice of the third dimension. (We also need the order of assets within each matrix to be the same, but that is already true.)

```

> aret.all <- cbind(crash=aret.crash, medium=aret.medium,
+   prosper=aret.prosper)
> aret.all
      crash medium prosper
equities   -20     8   25.0
bonds       6     4    3.5
commodities -9     5    8.0
cash        3     3    3.0

```

Min-Max Solution

The most common approach to simultaneous optimization is to maximize the minimum utility—also known as Pareto optimality. To clarify: each allocation will produce a utility for each scenario, the optimizer only pays attention to the worst utility for an allocation (with no regard to which scenario produces that utility); it then finds the allocation that does best with respect to the worst utility. The aim, then, is to never do really badly.

Before we do the actual optimization, we need to do some ground work.

```

> op.0 <- trade.optimizer(aprice, avar.all, aret.all,
+   gross.value=ten.thou, long.only=TRUE,
+   utility="mean-variance", iterations=0)
# (warning message omitted)
> utab3 <- op.0$utable
> utab3
      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9]
alpha.spot      0  1  2  0  1  2  0  1  2
variance.spot   0  0  0  1  1  1  2  2  2
destination     0  1  2  3  4  5  6  7  8
opt.utility     0  0  0  0  0  0  0  0  0
risk.aversion    1  1  1  1  1  1  1  1  1
wt.in.destination 1  1  1  1  1  1  1  1  1

```

The default behavior is to get all of the combinations of expected returns and variances—in this case we have three of each so there are 9 combinations. For our scenario analysis, though, we want only the combinations where the expected returns and the variances match—three combinations. In order to get the behavior that we want, we need to provide a suitable matrix as the `utable` argument. We could build a matrix from scratch, but it is slightly easier to start with the one that is wrong.

There are three changes that we need to make to this matrix. We need to select the three columns that we want, we need to reset the destinations so that they are numbered from zero through one less than the number of combinations, and we should specify the risk aversion parameter that we want.

```

> utab3 <- utab3[, c(1, 5, 9)]
> utab3["destination", ] <- 0:2
> utab3["risk.aversion", ] <- 0.02
> utab3

```

```

                                [,1] [,2] [,3]
alpha.spot                      0.00 1.00 2.00
variance.spot                    0.00 1.00 2.00
destination                      0.00 1.00 2.00
opt.utility                      0.00 0.00 0.00
risk.aversion                    0.02 0.02 0.02
wt.in.destination                1.00 1.00 1.00

```

At this point we are ready to do the actual optimization. In addition to needing the `utable` argument, the `quantile` argument needs to be 1 in order to get a min-max solution:

```

> op.minmax <- trade.optimizer(aprice, avar.all, aret.all,
+   gross.value=ten.thou, long.only=TRUE,
+   utility="mean-variance", utable=utab3, quantile=1)
> op.minmax$new.portfolio / 100
equities  bonds  cash
   0.91   34.82  64.27

```

A look at the utility values shows typical behavior of a min-max solution—the solution has the same utility for two of the scenarios.

General Simultaneous Optimization

Min-max optimization protects against the very worst outcome. However, it may really be buying too much insurance. We can change the `quantile` argument to a number that is less than 1 (but only numbers that are at least one-half are sensible):

```

> op.q7 <- trade.optimizer(aprice, avar.all, aret.all,
+   gross.value=ten.thou, long.only=TRUE,
+   utility="mean-variance", utable=utab3,
+   quantile=.7)
> op.q7$new.portfolio / 100
equities  bonds  cash
   0.76   32.70  66.54

```

In this case (of only three utility values) the quantity it is minimizing is a weighted average of the two worst utilities. For this example there is very little change from the min-max solution (and it seems to go more conservative rather than less conservative).

We can adjust the risk aversion but when `utable` is given, we need to use the `force.risk.aver` control argument. By default the risk aversion in the utility table that is passed in is taken as what is wanted.

```

> op.q7r1 <- trade.optimizer(aprice, avar.all, aret.all,
+   gross.value=ten.thou, long.only=TRUE,
+   utility="mean-variance", utable=utab3, quantile=.7,
+   risk.aversion=.01, force.risk.aver=TRUE)
> op.q7r1$new.portfolio / 100
equities  bonds  cash
   1.52   65.41  33.07

```


14.8 Compact Variance Objects

Usually the full variance matrix is given. In the case of multiple variances, a three-dimensional array is given where the third dimension represents the different variance matrices. Computationally, full matrices are the fastest, so the preferred format is full matrices as long as the memory of the machine on which the optimization is being done is large enough. This will almost certainly be the case unless the universe of assets is very large, or there is a large number of variance matrices.

caution

Unlike when a full matrix is given, there is no mechanism when giving compact variances for the optimizer to ensure that the assets are in the correct order within the variance. The user needs to take full responsibility that the asset order in the variances is the same as in the `prices` argument. This includes benchmarks, which enjoy some automatic treatment when full matrices are given.

In addition to full matrices, there are three formats supported:

- *simple factor models*. These have formula $\Lambda\Lambda^T + \Psi$ where Λ is a matrix of loadings with dimensions that are the number of assets by the number of factors, and Ψ is a diagonal matrix containing the specific variances. This will most likely be the result of a statistical factor model.
- *full factor models*. These have formula $\Lambda\Phi\Lambda^T + \Psi$ where Λ is a matrix of loadings with dimensions that are the number of assets by the number of factors, Φ is the variance matrix of the factors among themselves, and Ψ is a diagonal matrix containing the specific variances.
- *vech*. This is the lower triangle of the variance matrix stacked by column. So there is the variance of the first asset and the covariances of the first asset with all of the other assets, then the variance of the second asset and the covariances of the second asset with assets 3 and onward, etc.

The Variance List

When at least one variance is not full, the `variance` argument to the optimizers needs to be a list with special components. The components of the list are:

- `vardata`: Always required. This is a vector of the actual numbers for the variance representations all concatenated together. The order of the data in the representations is:
 - *full matrix*: The first column, the second column, etc. (This is, of course, the same as the first row, the second row, etc.)
 - *simple factor model*: The loadings for the first asset, the loadings for the second asset, etc. Then the specific variances.
 - *full factor model*: Each column of the factor variance matrix in turn. Then the loadings for the first asset, the loadings for the second asset, etc. Then the specific variances.

- *vech*: The vech representation.
- **vartype**: Always required. This is a vector of integers with length equal to the number of variance matrices represented. The integers indicate the type of representation for each matrix—there is no restriction on the combinations of types that may be used. The codes for the formats are:
 - *full matrix*: 0
 - *simple factor model*: 1
 - *full factor model*: 2
 - *vech*: 3
- **nvarfactors**: Required if any representation is a factor model (either simple or full). When given, this must have length equal to the length of **vartype**. The elements of this vector that correspond to a factor model must contain the number of factors for that model.
- **varoffset**: Never required, but may add some safety. This is an integer vector containing the offset for the start of each variance matrix within the **vardata** component. The first element is zero, the second is the length of the data for the first variance matrix, the third is the combined lengths of the first and second variance matrices, etc. The length of this vector is the number of matrices given, so the element that would be the total length of **vardata** is not given.

Here is an example of putting an object produced by `factor.model.stat` (from the `BurStFin` package) into this format:

```
> varfac <- factor.model.stat(retmat, out="fact")
> vdata <- c(t(varfac$sdev * varfac$loadings),
+   varfac$uniqueness * varfac$sdev^2)
> vlist <- list(vardata=vdata, vartype=1,
+   nvarfactors=ncol(varfac$loadings))
```

The `vdata` object is a vector of the pertinent numbers in the correct order. Note that the loadings are first multiplied by the standard deviations, then this matrix is transposed. Likewise the uniquenesses need to be scaled by the squared standard deviations. Since this is producing a simple factor model, the type of variance is 1. Finally, the number of factors is given as the number of columns of the loadings matrix. The `vlist` object is ready to be passed in as the **variance** argument.

caution

Be careful when annualizing a factor model representation. You want each element of the actual variance to be multiplied by some number—such as 252—but not every number in the representation is multiplied by the same thing.

Chapter 15

Dregs

This chapter contains topics that don't seem to belong anywhere else.

15.1 Portfolio Probe Constituents

The `pprobe.verify` function is useful for seeing if the installation on a particular machine worked okay. It is far from a full test suite—it merely checks that all of the functions are present, and that one particular problem optimizes okay. This can also be used to see which version of Portfolio Probe is present.

15.2 The Objectives

The objective function of an optimization has a numeric value as its result, and a set of inputs that can be changed. The job of the optimizer is to find the combination of inputs that achieves the best result for the objective function. By convention optimizers find the minimum of objective functions. The optimizer in `trade.optimizer` follows the convention.

The objective in `trade.optimizer` is the sum of two quantities: the negative of the utility (except “minimum variance” and “minimum distance” don't take the negative), and the penalty for broken constraints.. The penalty for a broken constraint is the appropriate element of `penalty.constraint` times a measure of how broken the constraint is. The penalties for all of the constraints are summed to get *the* penalty.

The same process is done for generating random portfolios. The difference is that the utility part is zero when `random.portfolio` is used—so there is only the penalty part that is minimized.

15.3 Writing C or C++ Code

The C code that underlies the optimization and random portfolio generation can be incorporated into your own C or C++ programs. There are quite a few arguments to the C functions for optimization, and some of the arguments involve a number of variables. Thus it is far from trivial to write the program calling these functions.

Speed of execution is not a good reason to go completely to C—the execution of the S code takes a fraction of a second and so is negligible for practical problems. Perhaps the most likely reasonable reason to stay all in C is to do real-time optimization within a C or C++ environment. However, doing real-time optimization in S is quite easy as shown on page 106.

If you are determined, there is a function in the package that can help you write C code. The `Cfrag.list` function will write a file that has the C initialization of variables given in an S language list. The list that you would want to do this for is an intermediate result of an optimization. For example:

```
> op.clist <- trade.optimizer(prices, varian, ...,
+   intermediate="Clist")
> Cfrag.list(op.clist[-1:-3], file="test.c")
```

The `test.c` file now holds a fragment of C code initializing most of the variables that are passed into the C function that does the optimization. This can be of use in two ways—either to initialize some variables that will always be the same (or at least mostly the same in your application), or to test your implementation as you are writing it.

To create a C program, you essentially need to recreate the `trade.optimizer` function. The S language code for `trade.optimizer` is the documentation for how to write the C code.

15.4 Bug Reporting

Send email to patrick@burns-stat.com to report any bugs that you find. A report should include:

- The operating system (including the version) that you are using.
- The version of R or S-PLUS that you are using.
- Give the `version` component (optimization) or attribute (random portfolios). For example:

```
> op$version
           C.code           S.code
"portgen_BurSt 1.17" "trade.optimizer 023"

> attr(randport, "version")
           C.code           S.code
"randport_BurSt 1.17" "random.portfolio 010"
```

- A full description of the problem—what happens, and what you wanted to happen.
- If S creates an error, then give the results of the call:

```
> traceback()
```

This displays the state that S was in when the error occurred.

- If possible, send the data for a small problem that reproduces the bug.
- If the bug is sporadic and uses `random.portfolio`, then also give the random seed from an object where the bug is observed:

```
> attr(ranport, "seed")
```

A reward is given to the first to report a bug in the code or documentation. Ideas for improving the ease of use or functionality are always welcome.

Bibliography

- [Almgren et al., 2005] Almgren, R., Thum, C., Hauptmann, E., and Li, H. (2005). Equity market impact. *Risk*.
- [Burns, 1998] Burns, P. (1998). *S Poetry*. <http://www.burns-stat.com>.
- [Burns, 2003a] Burns, P. (2003a). On using statistical factor models in optimizing long-only portfolios. Working paper, Burns Statistics, <http://www.burns-stat.com/>.
- [Burns, 2003b] Burns, P. (2003b). Portfolio sharpening. Working paper, Burns Statistics, <http://www.burns-stat.com/>.
- [Burns, 2009] Burns, P. (2009). The R Inferno. Tutorial, Burns Statistics, <http://www.burns-stat.com/>.
- [Burns et al., 1998] Burns, P., Engle, R., and Mezrich, J. (1998). Correlations and volatilities of asynchronous data. *The Journal of Derivatives*, 5(4).
- [diBartolomeo, 2003] diBartolomeo, D. (2003). Portfolio management under taxes. In Satchell, S. and Scowcroft, A., editors, *Advances in Portfolio Construction and Implementation*. Butterworth–Heinemann.
- [Grinold and Kahn, 2000] Grinold, R. C. and Kahn, R. N. (2000). *Active Portfolio Management*. McGraw–Hill.
- [Gupton et al., 1997] Gupton, G. M., Finger, C. C., and Bhatia, M. (1997). *CreditMetricsTM—Technical Document*. <http://www.riskmetrics.com>.
- [Kallberg and Ziemba, 1983] Kallberg, J. G. and Ziemba, W. T. (1983). Comparison of alternative utility functions in portfolio selection problems. *Management Science*, 29:1257–1276.
- [Michaud, 1998] Michaud, R. O. (1998). *Efficient Asset Management*. Harvard Business School Press.
- [Ulrich, 2009] Ulrich, J. (2009). *TTR: Technical Trading Rules*. <http://CRAN.R-project.org/package=TTR>. R package version 0.20-1.

Index

- 120/20 portfolio, 24, 48
- abbreviation of function arguments, 34
- abline function, 90
- allowance argument, 47, 48
- alpha table, 144–145
- alpha.constraint argument, 70
- alphas, *see* expected returns
- amortize cost, 119
- annualization
 - information ratio, 94
 - risk aversion, 140, 141
- argument abbreviation, 34
- array
 - three-dimensional, 153, 158, 161
- as.matrix function, 109, 110
- asset allocation, 99
- assets used in optimization, 91, 101
- asynchronous data, 111, 123
- atable, 149
- atable argument, 144
- backtest, 27
- bench.constraint argument, 71–72, 94, 95, 153
- benchmark
 - add to variance, 112
 - alpha table, 144–145
 - constraint, 94–95, 122, 153
 - dual, 95, 150–152
 - variance table, 144–145
- benchmark for long-short, 131–132
- benchmark-relative utility, 152
- bug report, 164
- build.constraints function, 58
- BurStFin package, 110
- buy-hold-sell list, 96–97, 103
- c function, 74, 76, 94
- C or C++ code, 163–164
- call component, 136
- cash flow, 98–99, 105–106
- CAUTION, 39, 49, 51, 111, 113, 116–118, 161, 162
- cbind function, 61, 129
- Cfrag.list function, 164
- cheatsheet, 126–127
- class function, 41
- close.number argument, 75
- comma-separated file, 40, 109, 114
- compact variance representation, 161–162
- compare to another program, 134
- constraint
 - absolute linear, 63–64
 - all, 45
 - benchmark, 94–95, 122, 153
 - building linear, 58–63
 - close number, 75
 - cost, 75
 - count, 67–69
 - country, 34
 - distance, 72–74
 - expected return, 69, 147–148
 - linear, 57–67
 - liquidity, 50–51, 55
 - long-side linear, 64
 - market capitalization, 62
 - monetary, 45–48
 - numerical linear, 61
 - penalty for breaking, 78–79, 113, 163
 - portfolio size, 93, 102
 - portfolio value, 106
 - quadratic, 76–78, 120
 - sector, 34
 - short-side linear, 64
 - soft, 78–79
 - sum of largest weights, 74–75
 - threshold, 52–53, 56

- tracking error, 34, 71–72, 94–95, 153
- trade, 63–64
- trade number, 93
- turnover, 93, 95, 102
- unsatisfiable, 41–42
- utility, 42
- variance, 70, 103, 145–147, 154
- violation, 65–66
- weight, 49–75
- constraint bounds, testing, 21–22, 84–86
- constraints.realized function, 65
- continuously compounded return, 111
- conventions, typography, 15
- Convergence, 125
- costs, 115
 - amortization, 119
 - asymmetric, 115
 - constraint, 75
 - linear, 116
 - nonlinear, 116–117
 - piecewise linear, 120
 - polynomial, 117–118
 - power law, 118–119
 - relative to expected returns, 119–120
 - relative to noise, 136
 - square root, 118–119
 - taxes, 118, 120
- count constraint, 67–69
- country constraint, 34
- cov.wt function, 110
- credit risk, 155–156
- csv file, 40, 109, 114
- currency, 121
- data frame, 109
- debugger function, 130
- debugging S, 129–130
- decay, time, 120
- deport function, 40
- deport functionexport function, 114
- dest.wt argument, 148, 149
- destination
 - of utility, 149
 - weight, 148, 149
- distance
 - constraint, 72–74
 - optimization, 97–98
- do.call function, 39
- do.warn argument, 117, 124–126
- dollar neutral portfolio, 24, 48
- doubleconst argument, 118
- drop function, 109, 110
- dual benchmarks, 95, 150–152
- duration, 61
- enforce.max.weight argument, 49
- error produced by S, 164
- examples, 92–99, 102–106, 156–160
- existing argument, 123
- expected returns, 93
 - multiple, 148, 150–160
 - relative to costs, 119–120
- factor model, 153
 - full, 161
 - fundamental, 153
 - macroeconomic, 153
 - simple, 161, 162
 - statistical, 110–112, 161
- factor vs factor, 111
- factor.model.stat function, 110, 162
- fail.iter argument, 42, 136
- feasible solution, 65
- file
 - read, 109
 - write, 40, 114
- fitted function, 112
- force.risk.aver argument, 150
- forced trade
 - automatic, 49
 - positions argument, 56
- forced.trade argument, 54, 104
- funeval.max argument, 134
- garbage, 121
- GARCH, 154
- gen.fail argument, 42
- generating random portfolios, 33–43
- genetic algorithm, 135
- global data, 111, 123
- gross value
 - definition, 48
- gross.value argument, 47
- head function, 38
- illiquid assets, 123
- import data into S, 109–110

- information ratio
 - annualizing, 94
 - definition, 93
 - maximize, 93–95, 102, 139
- init.fail argument, 41
- intermediate argument, 164
- intersect function, 61
- inventory model, 118
- iterations.max argument, 42, 134

- leverage, 106
- library function, 14
- limit.cost argument, 75
- lin.abs argument, 63
- lin.bounds argument, 58
- lin.constraint argument, 58
- lin.direction argument, 64
- lin.style argument, 59–60, 68
- lin.trade argument, 63
- linear constraint, 57–67
 - building, 58–63
- liquidity, 50–51, 55, 115
- log return, 111
- long and short lists, 103
- long value
 - definition, 47
- long-only portfolio, 91–100
- long-only with shorts, 131
- long-short portfolio, 101–108, 118
- long-side linear constraint, 64
- lot
 - round, 45
 - size, 114
- lower.trade argument, 50, 97, 104

- machine memory, 161
- market impact, 115
- market risk, 155
- max.weight argument, 49, 102, 128
- maximize information ratio, 93–95, 102, 139
- maximum expected return utility, 141
- maximum weight, 49
- mean-variance
 - optimization, 96, 104
- mean-variance utility, 140
- mean-volatility utility, 141
- median optimization, 148
- memory of the machine, 161

- min-max optimization, 148, 150, 159–160
- minimize tracking error, 92
- minimize variance, 92
- Minimum variance utility, 141
- miniter argument, 42
- missing values
 - variance estimation, 112
- monetary constraints, 45–48, 106
- multiple expected returns, 148, 150–160
- multiple time periods, 154
- multiple variances, 147, 150–161

- net asset value, 37
- net value
 - definition, 48
- nonlinear costs, 116–117
- not optimizing starting solution, 133–134
- ntrade argument, 51, 93, 97, 104

- objective function, 142, 163
- operating system, 164
- optimization
 - algorithm, 135
 - distance, 97–98
 - mean-variance, 96, 104
 - median, 148
 - min-max, 148, 150, 159–160
 - not doing, 133–134
 - objective, 142, 163
 - Pareto, *see* min-max optimization
 - real-time, 106–108
 - utility-free, 97–98
 - wrong, 121–124
- out.trade argument, 33, 34

- Pareto optimization, *see* min-max optimization
- passive portfolio, 92–93, 150–152
- penalty for broken constraint, 113, 163
- performance measurement, 19–21, 27, 83–84, 89–90
- polynomial costs, 117–118
- port.size argument, 52, 93, 95, 102
- portfolio
 - long-only, 91–100

- long-short, 101–108, 118
 - passive, 92–93, 150–152
- position
 - too small, 128
- positions argument, 54–127
- positions argument, 57
- ppoints function, 90
- pprobe.verify function, 163
- prices argument, 91, 101

- quadratic constraint, 76–78, 120
- quantile argument, 148, 153, 160
- quotes in S, 74, 154

- R language
 - instance of S language, 14
- R website, 13
- random portfolio, 33–43, 134
- random portfolio summary, 39
- random weights, 37
- random.portfolio function, 33–43, 45
- random.portfolio.utility function, 42
- randport.eval function, 38–39, 66
- rbind function, 129
- read a file, 109
- read.table function, 109
- real-time optimization, 106–108
- rep function, 97, 103
- repeat (S construct), 106
- return
 - continuously compounded, 111
 - log, 111
 - simple, 111
- rights issue, 122
- risk
 - credit, 155–156
 - market, 155
- risk aversion
 - parameter, 140, 150
- risk model, testing, 24, 86–89
- risk tolerance, 140
- risk.aversion argument, 150
- rival forecasts, 153–154
- round lot, 45

- S code note, 15, 33, 34, 39, 40, 61, 62, 74, 76, 97, 103, 106, 109, 111, 136, 153, 154, 157, 158
- S function
 - abline, 90
 - as.matrix, 109, 110
 - build.constraints, 58
 - c, 74, 76, 94
 - cbind, 61, 129
 - Cfrag.list, 164
 - class, 41
 - constraints.realized, 65
 - cov.wt, 110
 - debugger, 130
 - deport, 40, 114
 - do.call, 39
 - drop, 109, 110
 - factor.model.stat, 110, 162
 - fitted, 112
 - head, 38
 - intersect, 61
 - library, 14
 - ppoints, 90
 - pprobe.verify, 163
 - random.portfolio, 33–43, 45
 - random.portfolio.utility, 42
 - randport.eval, 38–39, 66
 - rbind, 129
 - read.table, 109
 - rep, 97, 103
 - summary, 39, 113
 - tail, 38
 - traceback, 130, 164
 - trade.optimizer, 91–108
 - update, 136
 - valuation, 114
 - var.add.benchmark, 112
- S language
 - debugging, 129–130
 - includes R, 14
 - multi-line command, 15
- S-PLUS website, 13
- scale.cost argument, 120
- scenario analysis, 156–160
- sector constraint, 34
- sense, making, 121
- short value
 - definition, 47
- short-side linear constraint, 64
- shrinkage, *see* Bayesian
- simple return, 111
- small portfolio, 128
- soft constraint, 78–79
- split, stock, 122

- stale prices, 123
- start
 - not optimizing, 133–134
- start.sol argument, 123, 134
- statistical factor model, 110–112, 161, 162
- stock split, 122
- stringency argument, 133, 136
- sum of largest weights constraint, 74–75
- summary function, 39, 113

- tab-separated file, 109, 114
- tail function, 38
- taxes, 118, 120
- three-dimensional array, 153, 158, 161
- threshold argument, 52–53, 127
- threshold constraint
 - positions argument, 56
- time decay, 120
- tol.positions argument, 57
- traceback function, 130, 164
- tracking error, 151
 - constraint, 34, 71–72, 94–95
 - minimize, 92
- trade
 - too trivial, 128
- trade.optimizer function, 91–108
- trading costs, *see* costs
- trading strategy, testing, 27–29, 90
- transaction costs, *see* costs
- troubleshooting, 127–129
- turnover, *see* trade.value argument
- turnover argument, 47, 93, 95, 102, 105
- txt file, 109, 114
- typography conventions, 15

- unchanging utility, 127
- universe.trade argument, 50, 56, 97, 104, 105
- update function, 136
- upper.trade argument, 50, 97, 104
- utable argument, 149, 160
- utility
 - benchmark-relative, 152
 - constraint, 42
 - definitions, 139
 - destination, 149
 - table, 149–150, 159–160
 - unchanging, 127
 - zero, 128
- utility-free optimization, 97–98

- valuation
 - function, 35–37, 114
 - too small, 128
- value constraints, 106
- var.add.benchmark function, 112
- var.constraint argument, 70
- variance, 110–112
 - adding benchmark, 112
 - argument, 161–162
 - compact representation, 161–162
 - constraint, 70, 103, 145–147, 154
 - factor model, *see* factor model
 - minimize, 92
 - missing value treatment, 112
 - multiple, 147, 150–161
 - table, 144–145, 149
 - vech representation, 161
- vech representation, 161
- version component, 164
- violation of constraints, 65–66
- volume, average daily, 50, 51, 118
- vtable argument, 144

- warnings, suppressing, 124–126
- weight
 - constraint, 49–75
 - destination, 148, 149
 - sum of largest constraint, 74–75
- write a file, 40, 114
- wrong optimization, 121–124

- zero utility, 128