4 Resampling and Portfolio Choice

Inputs for portfolio optimization problems are notorious for being measured with substantial estimation error. This is particularly troubling because optimization routines are often characterized as error maximization algorithms, leveraging errors in inputs rather than mitigating their effect. Consequently, financial economists and statisticians have relied on resampling techniques in order to understand the impact of estimation error in means and covariances (inputs) on the distribution of portfolio weights (outputs).[1] In statistics, resampling methods are referred to as bootstrap methods, and there are two basic types: the **parametric bootstrap**, where one fits a parametric model and samples from the fitted parametric model, and the **nonparametric bootstrap**, where one samples directly from the data without fitting a parametric model. See, for example, Efron and Tibshirani (1998) and Davison and Hinkley (1999) for details. In this chapter, we concentrate primarily on the parametric bootstrap using a fitted multivariate normal distribution, as is common in applications to finance.

Throughout the first three sections to follow, a simple numerical example will be used to illustrate the pitfalls of using the center of the resampled weight distribution for portfolio construction exercises. We need to rely on numerical examples in combination with Monte Carlo simulation, as no closed-form solutions are available.

4.1 Portfolio Resampling

Suppose we have estimated a mean vector and a covariance matrix of returns (in the following, we always assume returns come in the form of excess returns) from annual historical data with length n_{hist},

$$\hat{\mathbf{\Omega}}_0 = \begin{pmatrix} 400 & & \\ 210 & 255 & \\ 40 & 15 & 25 \end{pmatrix}, \quad \hat{\mu}_0 = \begin{pmatrix} 6.08 \\ 4.56 \\ 0.94 \end{pmatrix}, \quad n_{hist} = 30. \tag{4.1}$$

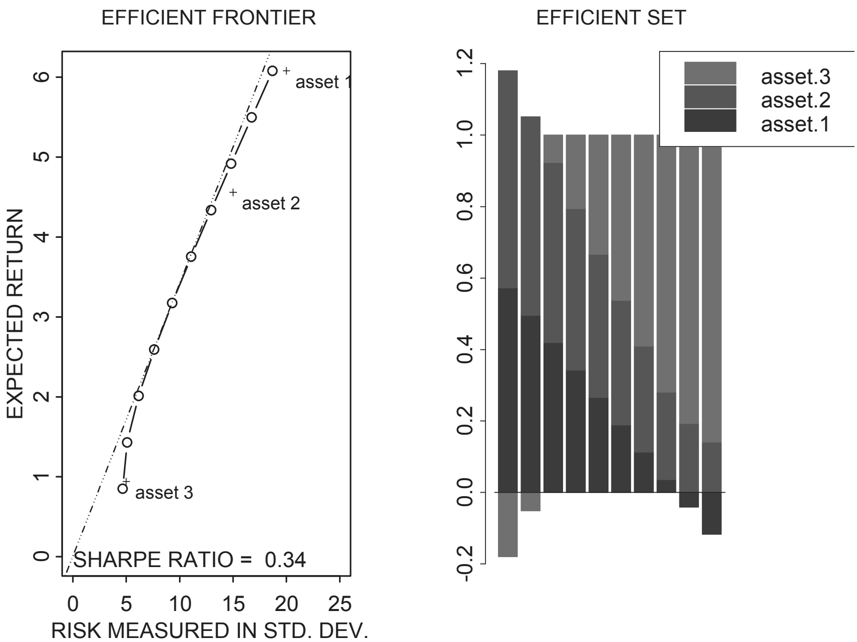

Figure 4.1 Markowitz Portfolios (with Short Selling Allowed)

The unconstrained efficient frontier and corresponding efficient set weights are shown in Figure 4.1. The maximum Sharpe ratio portfolio has weights

$$\mathbf{w}^*_{sharpe} = \frac{\hat{\mathbf{\Omega}}_0^{-1}\hat{\boldsymbol{\mu}}_0}{\mathbf{1}'\hat{\mathbf{\Omega}}_0^{-1}\hat{\boldsymbol{\mu}}_0} = \begin{pmatrix} \frac{1}{6} & \frac{2}{6} & \frac{3}{6} \end{pmatrix}^T .$$

Investors holding 100% in the maximum Sharpe ratio portfolio exhibit a risk aversion of $\lambda = {}^{\mu}\!/_{\sigma^2} = \mathbf{1}'\hat{\mathbf{\Omega}}_0^{-1}\hat{\boldsymbol{\mu}}_0 = 0.038$. As the maximum Sharpe ratio portfolio is the most prominent in finance, we will focus on this portfolio. With the exception of the minimum variance portfolio (which does not require return estimates), everything said in this chapter also applies to all other portfolios on the efficient frontier.

We know that $\hat{\mathbf{\Omega}}_0$ and $\hat{\boldsymbol{\mu}}_0$ have been estimated with error. In general, $\hat{\mathbf{\Omega}}_0$ is an $n \times n$ matrix, where n denotes the number of assets (here $n = 3$), whereas $\hat{\boldsymbol{\mu}}_0$ is an $n \times 1$ vector. The process of resampling will draw data for a number n_{draw} of returns for each of the n assets from the multivariate normal $N\left(\hat{\boldsymbol{\mu}}_0, \hat{\mathbf{\Omega}}_0\right)$. We can use the newly created block of data in the form of an $n_{draw} \times n$ matrix of asset returns (here 30×3) to construct a new mean vector and covariance matrix estimates $\hat{\mathbf{\Omega}}_1$ and $\hat{\boldsymbol{\mu}}_1$. It is often natural to set $n_{draw} = n_{hist}$, but this is not necessary. Obviously, the original and the resampled matrices will differ due to sampling error. The degree of difference will depend on n_{draw}. If we make n_{draw} small, our estimates will fluctuate greatly, while we will find much less difference for a large n_{draw}.[2] Repeating this procedure n_{sim} times, we create a large number of varying input vectors: $(\hat{\mathbf{\Omega}}_1, \hat{\boldsymbol{\mu}}_1, \ldots, \hat{\mathbf{\Omega}}_{n_{sim}}, \hat{\boldsymbol{\mu}}_{n_{sim}})$. We now ask ourselves what choices we would make if we repeatedly constructed optimal portfolios $\mathbf{w}_i$ from these resampled inputs and what insights can be gained from this exercise.

In order to ensure that decisions are indeed comparable across simulations, we assume that investors maximize $U(\mathbf{w}) = \mathbf{w}^T\boldsymbol{\mu} - \frac{\lambda}{2}\mathbf{w}^T\mathbf{\Omega}\mathbf{w}$, where the first-order conditions lead to the familiar formulas $\mathbf{w} = \lambda^{-1}\mathbf{\Omega}^{-1}\boldsymbol{\mu}$ $\boldsymbol{\mu} = \lambda\mathbf{\Omega}\mathbf{w}$ for optimal weights and implied returns. Note that in the current example $\lambda = 0.038$ remains constant through all simulations. As a start, we sample $n_{draw} = n_{hist} = 30$ returns from $N\left(\hat{\boldsymbol{\mu}}_0, \hat{\mathbf{\Omega}}_0\right)$ and compute the sample mean vector (for simplicity, covariances are assumed to be measured without error) and the corresponding optimal portfolio with a full investment constraint (i.e., weights need to add up to 100%). This is then repeated n_{sim} times for $n_{sim} = 1, 2, \cdots, 500)$.

We measure the distance between the center of the weight distribution and the original maximum Sharpe ratio portfolio that was constructed without taking estimation error into account (i.e., the maximum Sharpe ratio for the portfolio based on $\hat{\boldsymbol{\mu}}_0, \hat{\mathbf{\Omega}}_0)$ as the squared Euclidean distance

$$\left(\bar{\mathbf{w}} - \mathbf{w}^*_{Sharpe}\right)^T \left(\bar{\mathbf{w}} - \mathbf{w}^*_{Sharpe}\right), \text{ where } \bar{\mathbf{w}} = \frac{1}{n_{sim}}\sum_{i=1}^{n_{sim}} \mathbf{w}_i,$$

where $\mathbf{w}_i$ is the optimal weight vector for the i-th simulation.

It can be seen in Figure 4.2 that the distance between the center of the resampled distribution and the maximum Sharpe ratio portfolio converges to zero fairly rapidly as n_{sim} increases. Effectively this means that the center of the weight distribution recovers the original maximum Sharpe ratio portfolio. Alternatively, we can say that $\bar{\mathbf{w}} = \mathbf{w}^*_{Sharpe} + noise$, where the noise goes to zero fairly rapidly as n_{sim} increases. In this case, the use of resampling in creating new portfolios adds only noise to the portfolio construction.

Figure 4.3 visualizes the distribution of portfolio weights for $n_{sim} = 500$. Large positive or negative weights can occur in single simulation runs but will be averaged out. This is true for every number of draws n_{draw} per resampling.

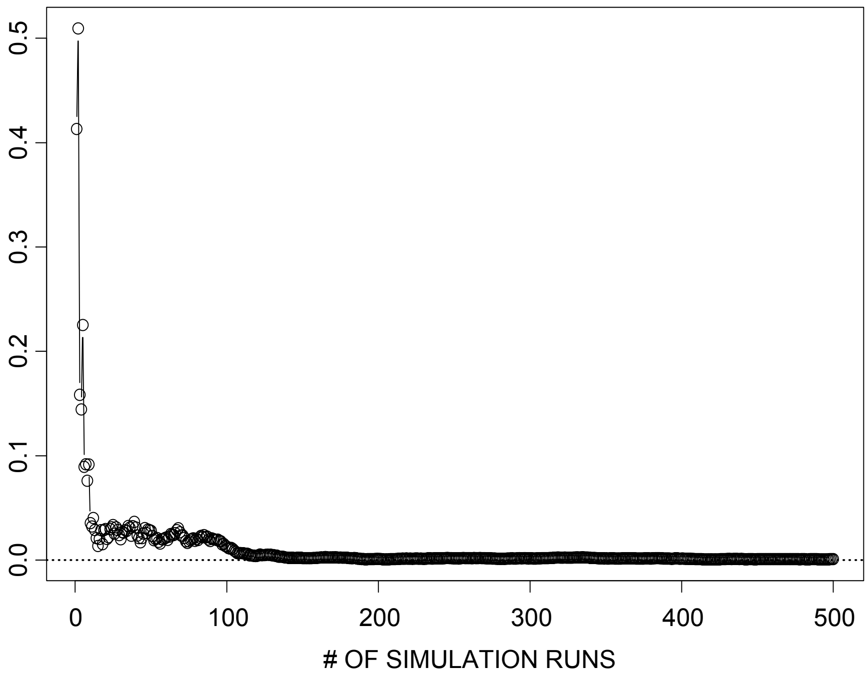

Figure 4.2 Resampling and Convergence (Short-Selling Allowed)

It is apparent from the results that repeatedly drawing average returns and subsequently averaging across optimally constructed portfolio weights, yields the same result as averaging across returns in the first place and then using the averaged returns for portfolio optimization. We could have seen that without having to go through the simulation exercise[3]:

$$\begin{aligned} \overline{\mathbf{w}} &= \frac{1}{n_{sim}} \sum_{i=1}^{n_{sim}} \mathbf{w}_i = \frac{1}{n_{sim}} \sum_{i=1}^{n_{sim}} \lambda^{-1} \mathbf{\Omega}^{-1} \hat{\boldsymbol{\mu}}_i \\ &= \lambda^{-1} \mathbf{\Omega}^{-1} \frac{1}{n_{sim}} \sum_{i=1}^{n_{sim}} \hat{\boldsymbol{\mu}}_i \\ &= \lambda^{-1} \mathbf{\Omega}^{-1} \overline{\boldsymbol{\mu}}. \end{aligned} \tag{4.2}$$

Note that we simulated the effect of estimation error on the distribution of portfolio weights. Neither the average portfolio nor its risk changed. However, we know that if investors are uncertain about their inputs, estimation risk will add to investment risk and the world will become a riskier place. Computing average weights based on resampling is unable to catch this effect, as it is not designed to do so.[4] However, straightforward bootstrap resampling of quantities such as the Sharpe ratio and the return and risk of the tangency portfolio can

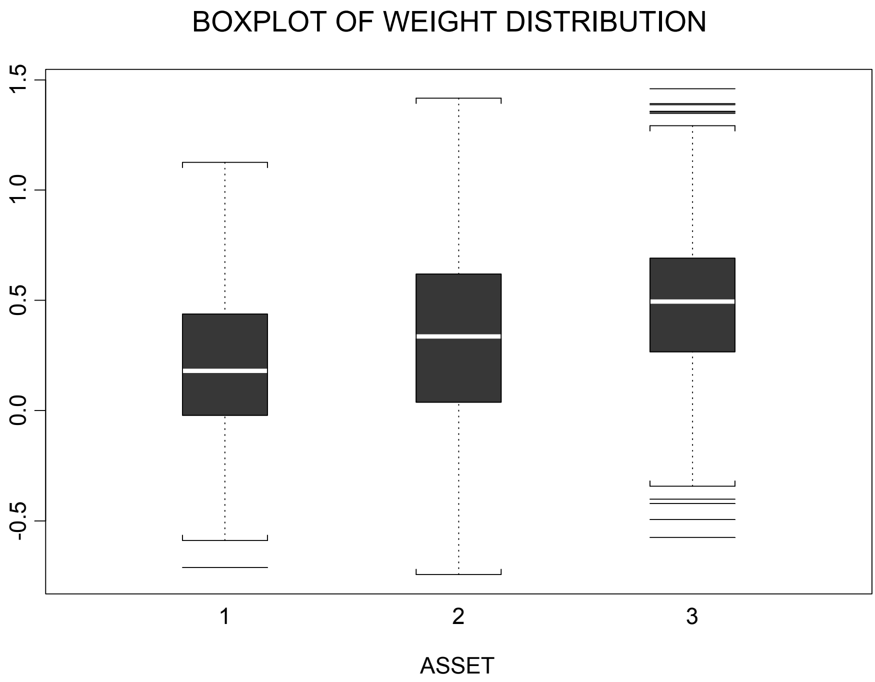

Figure 4.3 Distribution of Resampled Weights (Short-Selling Allowed)

indeed provide measures of uncertainty of operating points (see Sections 4.6 and 6.9.4).

In order to replicate the results above, readers can use Code 4.1, which works for both long/short `(short=T)` and long-only `(short=F)` optimization.

```
portfolio.resampling <- function(cov, fcst, n.sim,
  n.draw, short)
{
  resampled.pf <- matrix(0,ncol=ncol(cov),
     nrow=(n.sim+2))
  frontier.uc <- portfolioFrontier(cov,fcst,
     max.ret=max(fcst),n.ret=1000,
     unconstrained=short)
  iopt <- order(frontier.uc$returns/
     frontier.uc$sd)[1000]
  lambda <- frontier.uc$returns[iopt]/
     frontier.uc$sd[iopt]^2
  resampled.pf[(n.sim+1),] <-
     frontier.uc$weights[,1]
  resampled.pf[(n.sim+2),] <-
     frontier.uc$weights[,iopt]
```

```
  group <- matrix(rep(1, ncol(cov)), nrow=1)
  if(short==T) {
     bUP <- c(rep( Inf, ncol(cov)))
     bLO <- c(rep(-Inf, ncol(cov)))
  }
  if(short==F) {
     bUP <- c(rep(1, ncol(cov)))
     bLO <- c(rep(0, ncol(cov)))
  }
  cUP <- c(1)
  cLO <- c(1)
  for(i in 1:n.sim){
     x <- rmvnorm(n.draw, fcst, cov)
     cov.r <- var(x)
     fcst.r <- apply(x,2,mean)
     resampled.pf[i,] <- solveQP(-lambda*cov,
        fcst.r, group, cLO, cUP, bLO, bUP, ,
        type=maximize, trace=F)$variables$x$current
     cat(" run ", i, "\n")
  }
  list(resampled=resampled.pf)
}
```

Code 4.1 Portfolio Resampling and Weight Convergence

The first part of the function calculates the mean-variance frontier without estimation error. We can also infer the maximum Sharpe ratio portfolio from this (assuming expected returns and covariances are derived using the risk premium rather than total return).

4.2 Resampling Long-Only Portfolios

So far, we have allowed for short-selling in portfolio construction. We have seen that in this case the average resampled portfolio only adds noise to Markowitz portfolios. In this section, we will drop the possibility of going short in individual assets and return to more conventional portfolio optimization using a long-only constraint. Apart from this, we will perform the same calculations as in the previous section.

The first thing to note about Figure 4.4 is that distance (deviation from the estimation error-free solution) is much smaller when short-selling is not allowed, as the long-only constraint reduces the opportunities to leverage on information. We can also see that the distance measure in our simulations does not converge to zero. This means that repeatedly sampling with $n_{draw} = 30$ does

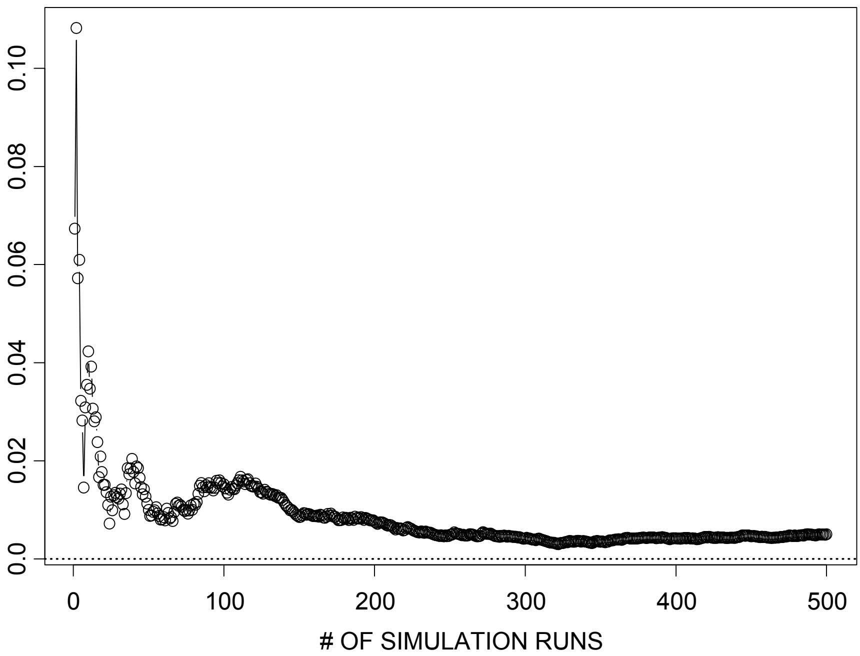

Figure 4.4 Resampling and Convergence (no Short-Selling Allowed)

not recover the Markowitz solution. Hence we get $\bar{\mathbf{w}} = \mathbf{w}^*_{Sharpe} + bias + noise$, where the noise goes to zero as the number of simulations increases but the bias does not. A look at Figure 4.5 provides the reason for this bias.

Weights that are less than zero due to a downward bias in some simulations can no longer be implemented. Hence, averaging will not lead back to the Markowitz solution, as individual assets are now either in or out but never short. The higher the volatility of an asset and/or the smaller n_{draw}, the more pronounced this effect will be. The next section will elaborate on this in more detail.

4.3 Introduction of a Special Lottery Ticket

In order to magnify the effect we just learned in the previous section and to show its relevance for asset allocation decisions, we will introduce a special lottery ticket with zero risk premium into our analysis. Lottery tickets are investments that offer diversification, as they are by definition uncorrelated with all other assets. Since our lottery ticket has zero expected excess return, it exposes investors to high volatility with no expected reward. Broadening the investment universe with lottery tickets should not improve the efficient frontier by pushing it up and to the left. Any asset allocation mechanism that systematically invests in lottery tickets should be treated with utmost caution. The following calculations are based on a lottery ticket with 60% volatility, 0% expected return, and zero covariance with existing assets.

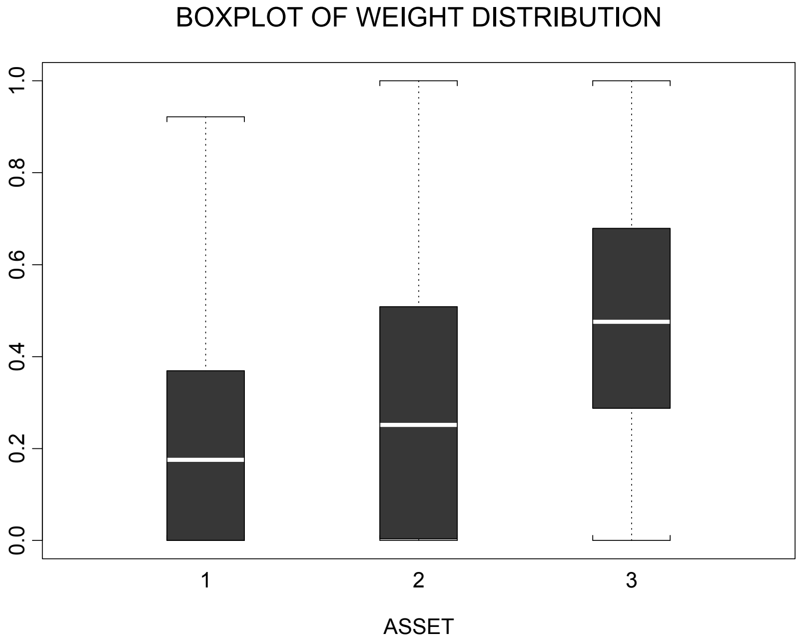

Figure 4.5 Distribution of Resampled Weights (no Short-Selling Allowed)

We repeat the previous calculations, where asset 4 represents the lottery ticket. Figure 4.6 and Figure 4.7 summarize the results. Note that the maximum Sharpe ratio portfolio derived from traditional mean-variance analysis does not allocate to the lottery ticket. Introducing a lottery ticket increases our distance measure in Figure 4.6 for a sufficiently large number of simulations. This should come as no surprise, as allocations to the lottery ticket amount to as high as 22% for some allocation runs, while we can never short the lottery ticket, even for those runs with large negative average returns. It is the long-only constraint that essentially transforms asset volatility into portfolio allocations. However, this does not necessarily mean that the higher the volatility of our lottery ticket the larger the allocation will become, as there are two separate effects at work. Higher volatility induces an upward bias into the average resampled weight, but at the same time higher volatility makes the lottery ticket less attractive, as it worsens the risk-return trade-off for any given risk aversion. While the first effect becomes obviously predominant for the maximum return portfolio, its exact trade-off depends on the risk aversion.

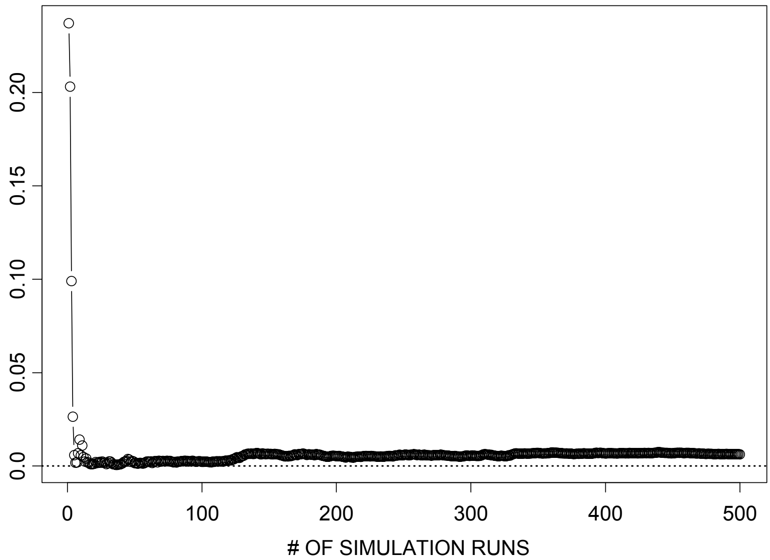

Figure 4.6 Resampling and Convergence (Lottery Ticket and Long-only Constraint)

Figure 4.8 shows that for a reasonably high risk aversion of $\lambda = 0.038$, increasing the volatility of the lottery asset will reduce the average allocation due to the higher risk. The volatility bias is still present, but the direct risk effect more than compensates for the upwards bias induced by high average returns for some simulation runs. For a low risk aversion of $\lambda = 0.01$, this effect also exists, but it starts at higher volatility levels. Up to a volatility level of 30%, the resampling bias dominates. From then on, the direct risk effect leads to smaller allocations even though the long-only constraint leads to more and more serious artifacts.

At this point, it is interesting to see what happens in a world that is affected by the same uncertainty about the correct inputs but that differs in institutional constraints. In short: does the introduction of a lottery ticket also have biased allocations if we are allowed to engage in short-selling? Note that allowing short-selling will not decrease the amount of estimation risk in the world. If anything, the opportunity to go short will increase the estimation error, as the optimizer can now establish long and short positions between similar highly correlated assets that look almost risk-free but yield large returns. Obviously, those almost arbitrage situations are most likely to be created by estimation error.

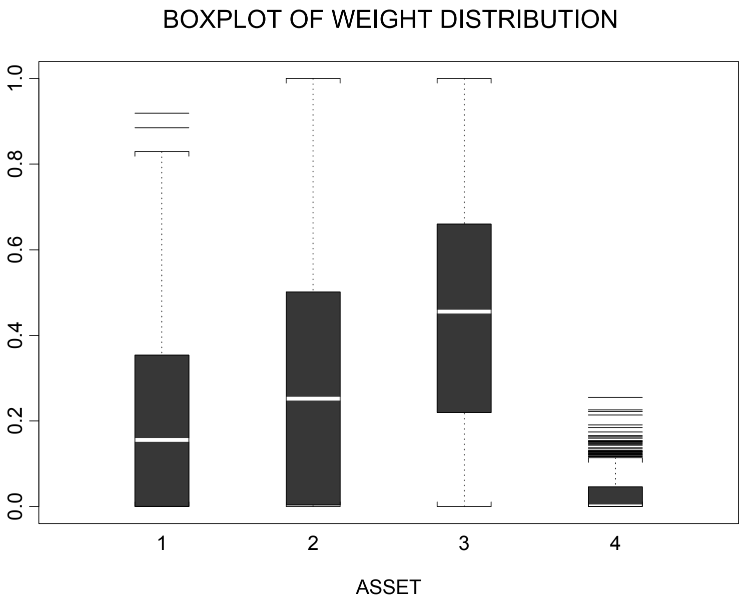

Figure 4.7 Distribution of Free Sampled Weights (Lottery Ticket and Long-Only Constraint)

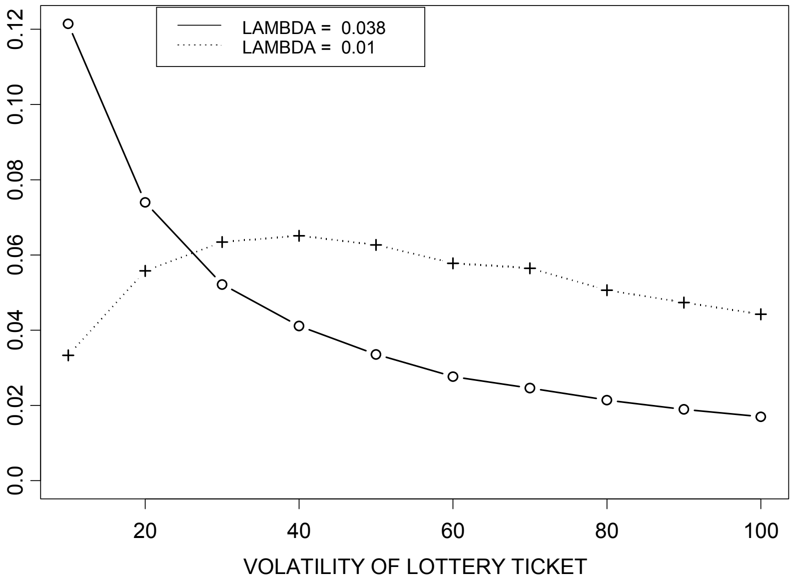

Figure 4.8 Risk Aversion and Volatility Bias

We see in the simulation results of Figure 4.9 that with short-selling allowed, the weights for the lottery ticket allocation scatter symmetrically around an average weight of 0%. Large positive allocations are counterbalanced on average by large negative allocations. Resampling without short-selling constraints helps us appreciate the dispersion in outcomes, while at the same time the average resampled weight is the same as the Markowitz weight.

Another way to look at portfolio resampling is to back out the implied returns of the average resampled portfolio. For $n_{draw} = 30$ and $n_{sim} = 500,$ we arrive at average resampled weights $\bar{\mathbf{w}} = \begin{pmatrix} 0.21 & 0.32 & 0.44 & 0.03 \end{pmatrix}^T$. In this case, one can check that the implied returns $\boldsymbol{\mu}_{implied} = \lambda \hat{\boldsymbol{\Omega}}_0 \bar{\mathbf{w}}$ differ substantially from our original forecasts $\boldsymbol{\mu}_{implied} = \begin{pmatrix} 6.45 & 4.67 & 0.92 & 4.15 \end{pmatrix}^T$. The risk premium for the lottery ticket in the latter case is more than 4%, compared with 0% for the portfolio with short-selling allowed. However, this is not plausible, as estimation error should not affect expected returns. By definition there is no uncertainty about expected returns. Estimation error without additional information should instead be reflected in the inflation of risk estimates, which now contain investment risk as well as estimation risk.[5]

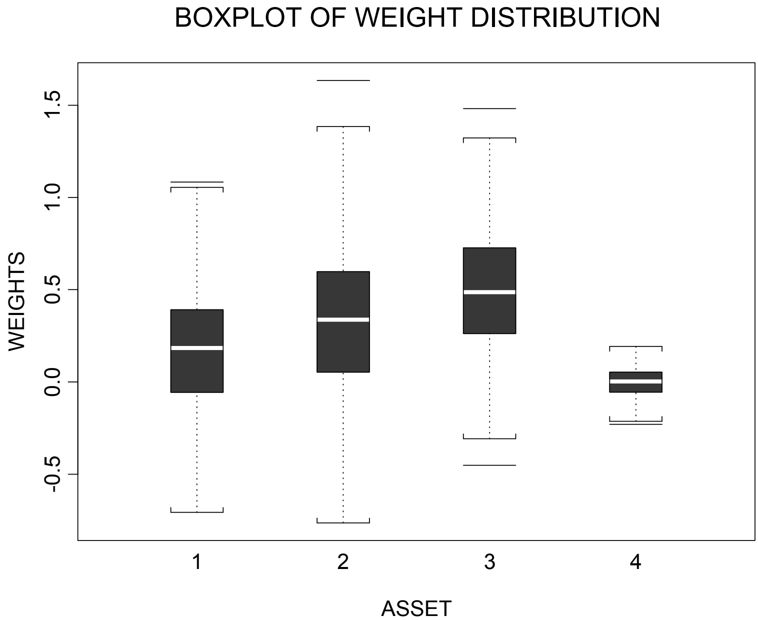

Figure 4.9 Distribution of Free Sampled Weights (Lottery Ticket without Long-only Constraint)

In order to appreciate the impact that the number of draws per resampling n_{draw} has on the allocation of our lottery ticket for a long-only portfolio, we repeat a large number of resamplings $n_{sim} = 100{,}000$ with antithetic variance reduction for various levels of n_{draw}. The results of these simulations are plotted in Figure

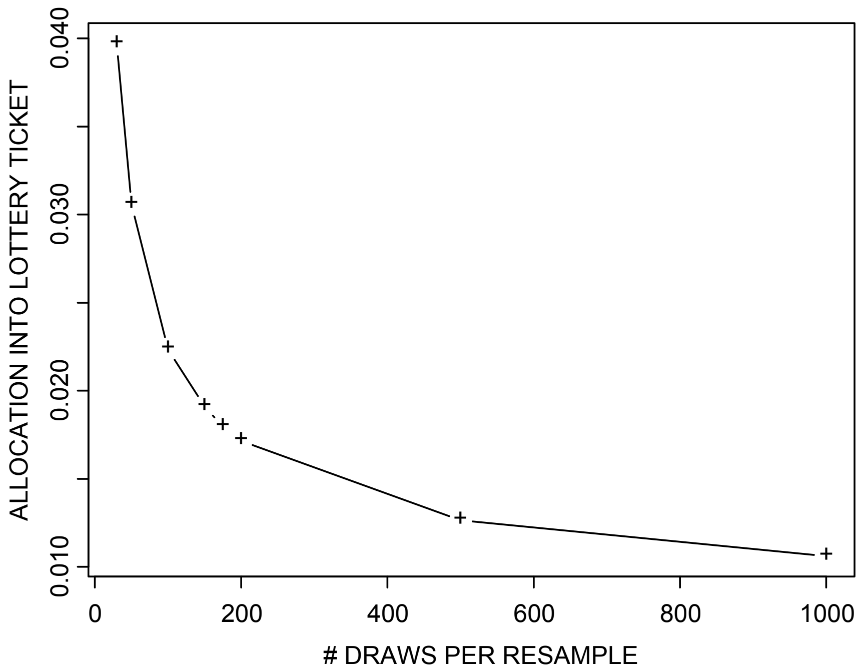

Figure 4.10 Allocation into Lottery Ticket versus n_{draw}

4.10. As the number of resamplings increases, and consequently the variance of estimated parameters decreases, the allocation into our lottery ticket decreases. At first sight this seems to be a confirmation of the concept of resampling. After all, a large number of draws per resampling means confidence in our inputs, in which case we would expect to recover the Markowitz solution. However, it is important to understand that no such effect exists if we allow short-selling. The average allocation into the lottery ticket would be independent of the number of draws even though the estimation error is the same. It is the long-only constraint that transforms asset volatility into asset allocation, implicitly raising the expected return for highly volatile assets.

4.4 Distribution of Portfolio Weights

The resampling procedure that results in a sequence of new covariance matrix and mean vector estimates allows us to generate a resampled set of optimal portfolio weights, thereby giving us an estimate of the distribution of portfolio weights. This in turn allows us to test whether two portfolios are statistically different using an appropriate distance in n-dimensional vector space. It may be tempting to use the simple Euclidean distance measure for the distance of a vector $\mathbf{w}_i$ of portfolio weights from the vector $\mathbf{w}_p$ of portfolio weights of another portfolio given by

$$(\mathbf{w}_p - \mathbf{w}_i)^T (\mathbf{w}_p - \mathbf{w}_i). \tag{4.3}$$

However, this is not the appropriate distance for correlated returns, and instead, under appropriate conditions, the proper statistical distance is given by

$$(\mathbf{w}_p - \mathbf{w}_i)^T \mathbf{\Omega}_{\mathbf{w}}^{-1} (\mathbf{w}_p - \mathbf{w}_i), \tag{4.4}$$

where $\bar{\mathbf{\Omega}}_{\mathbf{w}}$ is the variance-covariance matrix of portfolio weights $\mathbf{w}_i$ and $\mathbf{w}_p$ is the mean value of $\mathbf{w}_i$. When the $\mathbf{w}_i$ are normally distributed, this test statistic is distributed as a χ^2 with degrees of freedom equal to the number of assets. In the statistical literature this distance is known as the **Mahalanobis distance**, and an intuitive explanation of the distance is provided in Section 6.6.[6]

Suppose for simplicity that we have two assets with 10% mean and 20% volatility each. Suppose further that the correlation between the two assets is zero and the risk aversion coefficient is $\lambda = 5$. The optimal solution without estimation error is given below:

$$\mathbf{w}^* = \begin{bmatrix} w_1^* \\ w_2^* \end{bmatrix} = \lambda^{-1}\mathbf{\Omega}^{-1}\boldsymbol{\mu} = 0.2\begin{bmatrix} \frac{1}{(0.2)^2} & 0 \\ 0 & \frac{1}{(0.2)^2} \end{bmatrix}\begin{bmatrix} 0.1 \\ 0.1 \end{bmatrix} = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}. \tag{4.5}$$

A resampled version of (4.5) is now easily obtained with a few lines of S-PLUS code (see Code 4.2), assuming that the returns are normally distributed.

```
# inputs
Cov <- diag(rep(0.2^2,2))
mu.bar <- c(rep(0.1,2))
n.sim <- 1000
n.draw <- 60
lambda <- 0.2
# simple resampling function
resampling <- function(Cov, mu.bar, n.sim, n.draw,
  lambda)
{
  resampled.weights <- matrix(0, n.sim, ncol(Cov))


  for(i in 1:n.sim) {
     resampled.returns <- rmvnorm(n.draw, mu.bar,
```

```
        Cov)
      VarCov <- var(resampled.returns)
      Mean <- apply(resampled.returns,2,mean)
      w <- lambda*solve(VarCov)%*%Mean
      resampled.weights[i,] <- t(w)
   }
   list("resampled.weights"=resampled.weights)
}
# plot results
x <- resampling(Cov, mu.bar, n.sim, n.draw,
   lambda)$resampled.weights
plot(x[,1],x[,2],xlab="weight asset 1",
   ylab="weight asset 2")
```

Code 4.2 Portfolio Resampling and Weight Distribution

Note that for illustrative purposes we have calculated optimal portfolios without full investment constraints in the code above. Because these portfolios do not require holdings to add up to one, one might be tempted to conclude that these are not portfolios. But one could think of cash as a third (filling) asset, as cash would leave the marginal risks of the portfolio, as well as the total risk of the risky portion of the portfolio, unchanged. While the optimal solution weight is 50% for both assets, Figure 4.11 shows that the estimated weights are scattered around this solution. Comparing the vector difference with an appropriate percentage point (e.g., the upper 95% point) of a chi-squared distribution with two degrees of freedom yields a measure of how statistically different a portfolio is from the optimum.

From the definition of optimal weights, one sees that for our simple example the covariance matrix of the resampled weights is given by

$$\begin{aligned}\mathbf{\Omega_w} &= \operatorname{cov}(\lambda - 1\mathbf{\Omega}^{-1}\hat{\boldsymbol{\mu}}) \\ &= \lambda^{-2}\mathbf{\Omega}^{-1}\operatorname{cov}(\hat{\boldsymbol{\mu}})\mathbf{\Omega}^{-1} \\ &= \lambda^{-2}\mathbf{\Omega}^{-1}\frac{\mathbf{\Omega}}{n_{draw}}\mathbf{\Omega}^{-1} \\ &= \frac{(.2)^2}{n_{draw}}\mathbf{\Omega}^{-1} \\ &= \frac{1}{n_{draw}}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}.\end{aligned} \tag{4.6}$$

Since $n_{draw} = 60$ in our example, this gives $1/\sqrt{60} = .13$ as the standard errors of the weights. This is consistent with the display in Figure 4.11.

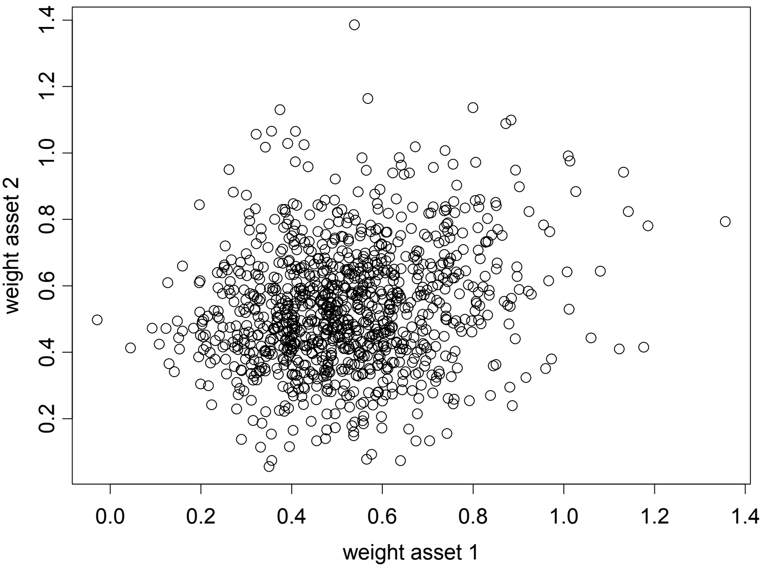

Figure 4.11 Estimation Error and Portfolio Weights

We remark that, for fully invested portfolios, the n-dimensional vector of weights will lie on an $n-1$ dimensional hyperplane that intersects the coordinate axes at the value one. (In this simple two-dimensional case, the weights lie along a line through the points (0,1) and (1,0).) In such cases, we can simply look at the distribution of $n-1$ of the weights in the $n-1$ dimensional subspace. In our simple example above, this would amount to looking at just one weight, which is not very interesting.

Now, to be ever so slightly realistic, let's consider the estimated bivariate distribution of the weights based on observed data that are assumed to be normally distributed according to an estimated mean and covariance for the weights, $\hat{\mathbf{w}}^*$ and $\widehat{\boldsymbol{\Omega}}_w$. For simplicity, we will use the true optimal weights $\mathbf{w}^* = (.5,.5)'$, leaving it to the reader to repeat the experiment with $\hat{\mathbf{w}}^*$, and compute the estimate $\widehat{\boldsymbol{\Omega}}_w$ directly from the resampled weights (rather than using the previous formula). The resulting bivariate density is

$$p(w_1, w_2) = \frac{1}{2\pi \det\left(\hat{\boldsymbol{\Omega}}_w\right)^{\frac{1}{2}}} e^{-\frac{1}{2}\begin{bmatrix} w_1 - w_1^* \\ w_2 - w_2^* \end{bmatrix}' \hat{\boldsymbol{\Omega}}_w^{-1} \begin{bmatrix} w_1 - w_1^* \\ w_2 - w_2^* \end{bmatrix}} .$$

Code 4.3 generates the perspective plot of Figure 4.12 and the contour plot of Figure 4.13. For the case depicted in these figures, the estimated inverse covariance matrix of the weights was

$$\hat{\mathbf{\Omega}}_w^{-1} = \begin{bmatrix} 27.93 & 0.005 \\ 0.005 & 27.76 \end{bmatrix}.$$

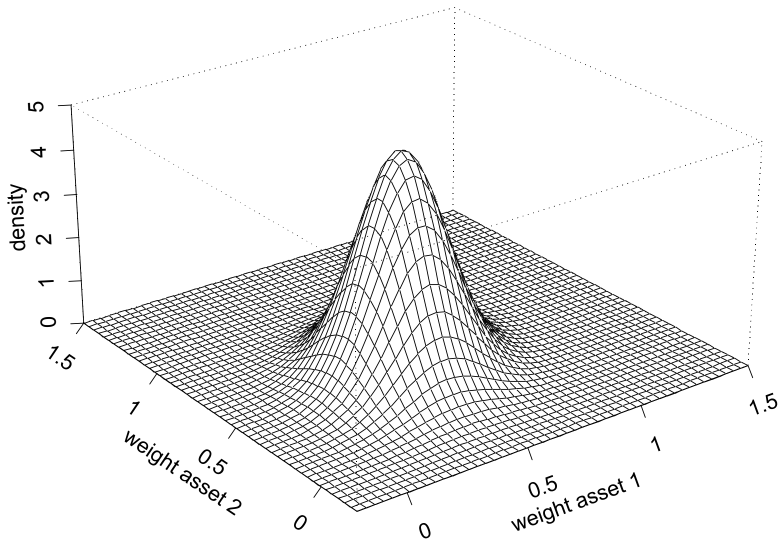

Figure 4.12 Bivariate Normal Weight Distribution for Resampled Portfolio Weights

```
Cov.w <- var(x)
w1 <- seq(-0.2, 1.5,length=100)
w2 <- seq(-0.2, 1.5,length=100)

f1 <- function(w1,w2)
{
  S <- solve(Cov.w)
  d <- (w1-0.5)^2*S[1,1]+(w2-0.5)^2*S[2,2]+
      2*(w1-0.5)*(w2-0.5)*S[1,2]
  1/(2*pi*sqrt(det(Cov.w)))*exp(-1/2*d)
}

z <- outer(w1,w2,f1)
graphsheet()
persp(w1, w2, z, xlab="weight asset 1",
```

```
  ylab="weight asset 2", zlab="density")
graphsheet()
contour(w1, w2, z, nlevels=10,
  xlab="weight asset 1", ylab="weight asset 2")
points(x[,1], x[,2])
```

Code 4.3 Portfolio Resampling and Weight Distribution

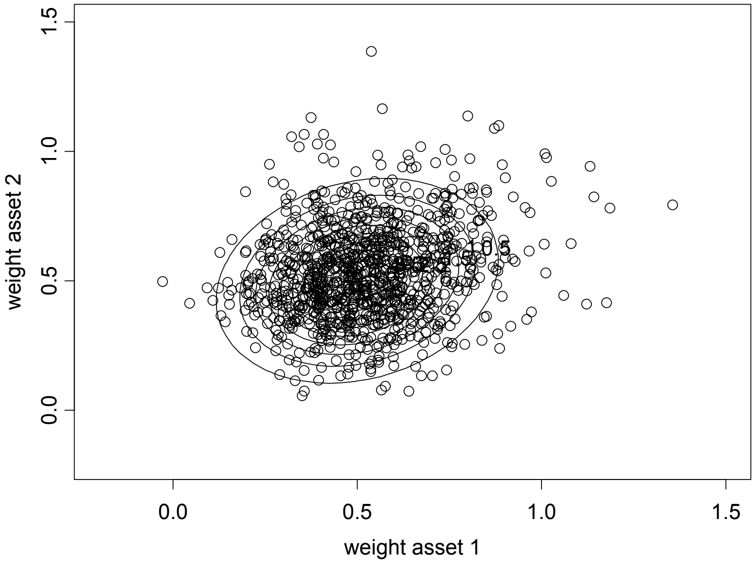

Figure 4.13 Weight Distribution and Lines of Constant Density

Michaud (1998) uses a different distance measure that is widely applied in asset management. His measure recognizes that two portfolios with the same risk and return might actually exhibit different allocations. The distance between two portfolios is defined as

$$\left(\mathbf{w}_p - \mathbf{w}_i\right)^T \hat{\mathbf{\Omega}}_0 \left(\mathbf{w}_p - \mathbf{w}_i\right), \tag{4.7}$$

which is equivalent to the squared tracking error. The procedure runs as follows:

Step 1. Define a portfolio against which to test the difference. Calculate (4.4) for all resampled portfolios.

Step 2. Sort the portfolios by tracking error in descending order (highest on top).

Step 3. Define TE_α as the critical tracking error for the $\alpha\%$ level (i.e., if 1000 portfolios are resampled and the critical level is 5%, then look at the tracking error of a portfolio that is 50th from the top). Hence, all portfolios for which $\left(\mathbf{w}_p - \mathbf{w}_i\right)^T \hat{\boldsymbol{\Omega}}_0 \left(\mathbf{w}_p - \mathbf{w}_i\right) \geq \mathrm{TE}_\alpha^2$ are labeled statistically different.

Step 4. Calculate the minimum and maximum allocations for each asset within the acceptance region.

For a three-asset example, the uncertainty about the optimal weights can be visualized, but it becomes “quite hard” for higher dimensions.

It should be noted that similarity is defined with regard to the optimal weight vector rather than in terms of risk and return. Two portfolios could be very similar in terms of risk and return but very different in allocation. This is well-known, as risk/return points below the frontier are not necessarily unique. Even so, this test procedure is intuitive. It should be noted that the dispersion in weights is large, so it will be difficult to reject the hypothesis that both portfolios are statistically equivalent even if they are not. The power of the suggested test is expected to be low.

4.5 Theoretical Deficiencies of Portfolio Construction via Resampling

4.5.1 Aggregation Problems

Constructing “optimal” portfolios using portfolio resampling requires that we average portfolios in some way (e.g., we average portfolios that carry either the same rank or the same risk-return trade-off).[7] In the case of no long-only constraints, the concept of resampled efficiency will coincide with Markowitz efficiency in the large sample limit (i.e., resampled efficiency in finite sample sizes equals Markowitz efficiency plus noise). Note that even though all inputs are measured with error, resampled efficiency will not pick this up. Asset risk remains unchanged even though the world becomes much riskier in the presence of estimation error.

In the case of long-only constraints, the situation changes considerably. As assets can never be short, we will see that for some resamplings the maximum return portfolio will be 100% cash. This leads to a sampling of cash into the maximum return portfolio. Another consequence is that we cannot engineer portfolios that exhibit low λ’s (without long-only constraints, we could have always shorted assets with a negative risk premium), which makes the similarity of rank- and lambda-based approaches questionable. Note also that the inclusion of cash in the maximum return portfolio contrasts both with intuition and portfolio theory. In the case of estimation error, investors will still hold a combination of cash and a market portfolio with the same composition as in the case of no estimation error, with more weight being put on cash, as cash carries no investment risk and is free of estimation error.

Finally, we note that the average is a poor indicator for the center of a distribution that is asymmetric due to heavy truncation at both ends (between 0% and 100%).

4.5.2 Overdiversification

A portfolio construction methodology that allocates to *every* single asset in the universe across *all* portfolios along the efficient frontier creates overdiversification. The combination of the long-only constraint and portfolio resampling will allocate even to dominated assets as long as a lucky draw makes them attractive, while the worst that can happen in all other allocations is a zero weight. Hence, the increase in risk per unit of expected return is not due to estimation error but rather due to overdiversification.

4.5.3 Optionality Problem

Suppose two assets possess the same expected return but one of them has a significantly higher volatility. One could think of this as an international fixed income allocation on a hedged and unhedged basis. Most practitioners (and the mean-variance optimizer) would exclude the higher-volatility asset from the solution unless it has some desirable correlations. How would resampled efficiency deal with these assets? Repeatedly drawing from the original distribution will result in draws for the volatile asset with highly negative returns as well as highly positive returns. Quadratic programming will heavily invest in this asset in the latter case and short the asset in the former case. However, as shorting is not allowed for portfolios with long-only constraints, this will result in positive allocation for draws with high positive average return and zero allocations for draws with high negative average return. This is different from an unconstrained optimization, where large long positions would be offset on average by large negative positions. Consequently, an increase in volatility will yield an increase in the average allocation, and a worsening Sharpe ratio would be accompanied by an increase in weight. This is not a plausible result. It arises directly from the averaging rule in combination with a long-only constraint that results in assets being either in or out but never negative. This behavior is a kind of optionality in which the holder is hurt in terms of bias in the weights whenever a long-only constraint forces otherwise negative coefficients to be positive and less than one in value.

4.5.4 Statistical Foundation Issues

Estimation Error Heritage. All resamplings are derived from the same initial estimates $\hat{\boldsymbol{\Omega}}_0$, $\hat{\boldsymbol{\mu}}_0$ of the covariance matrix and mean returns. However, the true distribution is unknown. Hence, all resampled portfolios will suffer from the deviation of the estimates $\hat{\boldsymbol{\Omega}}_0$, $\hat{\boldsymbol{\mu}}_0$ from their true values $\boldsymbol{\Omega}_{true}$, $\boldsymbol{\mu}_{true}$, in the same way. Averaging will not help very much in this case, as the averaged weights are the result of an input vector, which itself is very uncertain. Hence, it is fair to say that all resampled portfolios inherit the same fundamental estimation error. The utility of normal distribution parametric resampling relies on the assumption that $\hat{\boldsymbol{\Omega}}_0$, $\hat{\boldsymbol{\mu}}_0$, is reasonably close to $\boldsymbol{\Omega}_{true}$, $\boldsymbol{\mu}_{true}$. If this is not the case, the estimation error in $\hat{\boldsymbol{\Omega}}_0$, $\hat{\boldsymbol{\mu}}_0$ is passed on to $\hat{\boldsymbol{\mu}}_1, \hat{\boldsymbol{\Omega}}_1$, $\hat{\boldsymbol{\mu}}_2, \hat{\boldsymbol{\Omega}}_2$, ..., which one might call "estimation error heritage" (see Figure 4.14).[8]

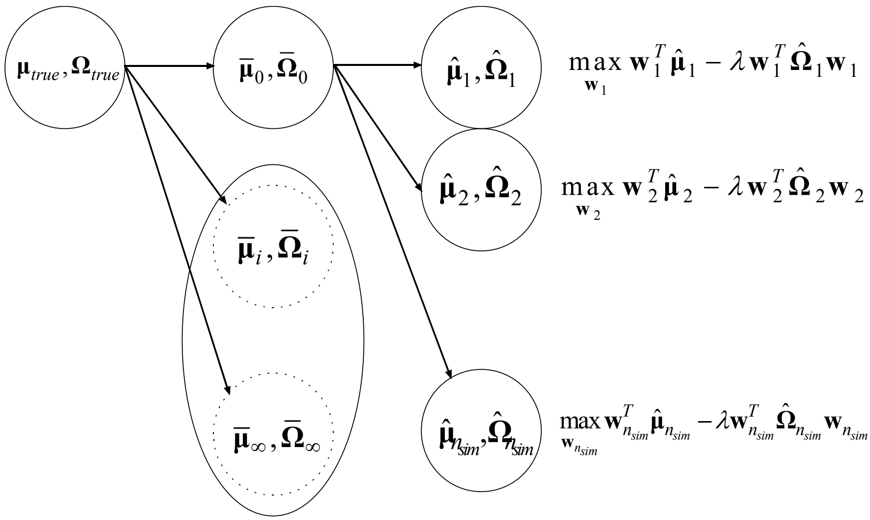

Figure 4.14 Resampling and Estimation Error Inheritance

Parametric Bootstrap Limitations. We know that asset returns are not normally distributed and in some cases are quite non-normal. That makes use of a normal distribution parametric bootstrap highly suspect. One might well turn to a multivariate non-normal distribution such as a multivariate t distribution. This requires careful estimation of the degrees of freedom and robust estimation of the mean vector and covariance matrix. Another approach is to use a nonparametric bootstrap that makes no assumptions about the distribution of the returns, as discussed in Sections 4.6 and 6.9.4.

Resampling Bayes. Sometimes it is argued that $\hat{\boldsymbol{\Omega}}_0$, $\hat{\boldsymbol{\mu}}_0$ does not need to be estimated from historical data but can also be the result of Bayesian calculations. However this is entirely against the spirit of Bayesian statistics. Once we calculate the predictive distribution, we have already put in all our subjectivity, and the Bayesian has to accept the result. Resampling from predictive distributions in order to construct better portfolios is pointless.

4.5.5 Lack of Decision-Theoretic Foundation

Resampled efficiency has no decision-theoretic foundation and as such it is questionable whether its use is fiduciary. What resampling actually achieves is some sort of return shrinkage in the presence of long-only constraints. Backing out implied returns from average resampled portfolios already revealed to us that low returns of relatively high-risk assets tend to be adjusted upward and vice versa. The advantage of this form of shrinkage over classical shrinkage methods is that portfolios constructed from it add up to 100%. This is not the case for the statistical shrinkage model, which in addition may still lead to concentrated corner portfolios. However, while we have perfect control over the latter, this cannot be said about the implied returns from resampling.

4.6 Bootstrap Estimation of Error in Risk-Return Ratios

4.6.1 The Problem

Reported risk-return ratios relate average returns to alternative measures of risk and hence involve the ratio of a random numerator and denominator (due to sampling error). As such, point estimates of these ratios are easy to calculate, but confidence intervals are much more difficult to obtain. However, we need confidence intervals for any kind of statistical inference (and hence for decision making). While asymptotic normal distributions have been obtained for the Sharpe ratio under idealized conditions,[9] the idealized conditions do not always hold, and furthermore asymptotic distributions may be poor approximations in finite sample sizes. There is little guidance on the small-sample behavior of risk-adjusted performance measures or on the number of data points needed to justify the use of asymptotic results. Moreover, these analytical solutions are either extremely difficult to work out or simply do not exist for modifications of the popular Sharpe ratio that focus more on downside risk. As an example, we look at the well-known Sortino ratio which relates average return to the standard deviation of downside returns. What we need is a general method that provides us with standard errors and confidence intervals for arbitrary risk-return ratios, sample sizes, and distributions.

4.6.2 Bootstrapping Theory as an Alternative

Suppose we observe a series of excess returns $r_1, r_2, \cdots, r_m$.[10] Ex-post-risk-return ratios $\hat{\varsigma}$ are calculated as the ratio of the average return per unit of risk. For illustrative purposes, we focus on the Sharpe and Sortino ratios given below. Both ratios differ with respect to the risk measure used. The sample calculations for these two ratios are

$$\begin{aligned} \textit{Sharpe ratio} &= \frac{\frac{1}{m}\sum_{i=1}^{m} r_i}{\sqrt{\frac{1}{m-1}\sum_{i=1}^{m}(r_i - \bar{r})^2}}, \\ \textit{Sortino ratio} &= \frac{\frac{1}{m}\sum_{i=1}^{m} r_i}{\sqrt{\frac{1}{m-1}\sum_{i=1}^{m} I(r_i < 0)(r_i)^2}}, \end{aligned} \tag{4.8}$$

where $I(r_i < 0)$ denotes the indicator function. The Sharpe ratio[11] employs the symmetric standard deviation of returns risk measure in the denominator, equally penalizing downside and upside deviations from the sample mean return. The denominator asymmetric risk measure in the Sortino ratio[12] includes only negative returns in its calculation of squared returns. High ratios are preferable, everything else being equal, as they indicate a better return per unit of risk taken.

We include the Sortino ratio for three reasons. First, it better captures the risk if returns are non-normally distributed, as is the case for hedge fund returns for example, and is particularly relevant when the distribution has a negative skew. Second, it is well-known that the Sortino ratio suffers more from estimation error, as it uses roughly half as many data points in the denominator risk measure relative to the Sharpe ratio. Third, no large sample approximations exist. If the small sample distribution of $\hat{\varsigma}$ is far from normal, classical methods are biased and unreliable.

In any case, the analytic formulas for the large sample distributions of the ratios above are extremely hard to come by. In order to overcome this problem we rely on nonparametric bootstrapping techniques. Nonparametric bootstrap resampling treats the empirical distribution function of the current sample as *the* nonparametric approximation of the true distribution—in the absence of further information, it is the best we have. It then repeatedly draws from the empirical distribution and recalculates the statistic of interest many times to arrive at the bootstrap sampling distribution. The bootstrap sampling distribution can then be used to construct standard errors, confidence intervals, and hypothesis tests; see, for example, Efron and Tibshirani (1998) and Davison and Hinkley (1999).

As an example, suppose that we are given 160 monthly returns on the HFR fund-of-funds index ranging from January 1990 to April 2003. We use the JPM one month cash rate from DataStream to calculate a risk-free rate. The nonparametric bootstrapping procedure is as follows.

1. Randomly draw 160 (original sample size) returns with replacement from the original sample.
2. Calculate a new risk-return ratio $\hat{\varsigma}_b^*$ based on the resampled returns.
3. Repeat this procedure for $b = 1,...,B$ times, arriving at $\hat{\varsigma}_1^*, \hat{\varsigma}_1^*, \cdots, \hat{\varsigma}_b^*, \cdots \hat{\varsigma}_B^*$ resampled ratios.

The bootstrap sampling distribution of $\hat{\varsigma}_b^*$ can now be used to judge whether the sampling distribution of $\hat{\varsigma}$ for small samples is normal and hence whether traditional sampling theory approximations might not be so bad after all. Setting $B = 10,000$ and using Code 4.4 along with the S-PLUS functions `qqplot` and `histogram`, we get the results in Figure 4.15 and Figure 4.16 . In Figure 4.15, we see that for the Sharpe ratio all resampled realizations plot very close to a straight line, and so we conclude that the Sharpe ratio is quite normally distributed. The same cannot be said about the Sortino ratio, for which the normal Q-Q plot has substantial deviations from linearity at both ends, being heavy-tailed to the right and short-tailed to the left. As we suspected, the histograms in Figure 4.16 show that the Sortino ratio has a much larger dispersion in resampled outcomes than the Sharpe ratio and hence a much larger estimation error. While a small-sample normal approximation looks reasonable for the traditional Sharpe ratio, such an approximation is likely to be largely misleading for the Sortino ratio.

```
B <- 10000
sharpe.ratio <- function(x){
  mean(x,na.rm=T)/stdev(x,na.rm=T)
}
sortino.ratio <- function(x){
  mean(x,na.rm=T)/sqrt(mean(pmin(x,0)^2,na.rm=T))
}
simple.bs <- bootstrap(x,sortino.ratio,B)
```

Code 4.4 Simple Bootstrap

We now use the 2.5% and 97.5% percentiles of the bootstrap distribution to obtain a symmetric 95% confidence interval $CI\left(\hat{\varsigma}_{2.5\%}^*, \hat{\varsigma}_{97.5\%}^*\right)$. The Sortino ratio confidence interval is (.11, .92), and the Sharpe ratio confidence interval is *CI*(.08, .41).

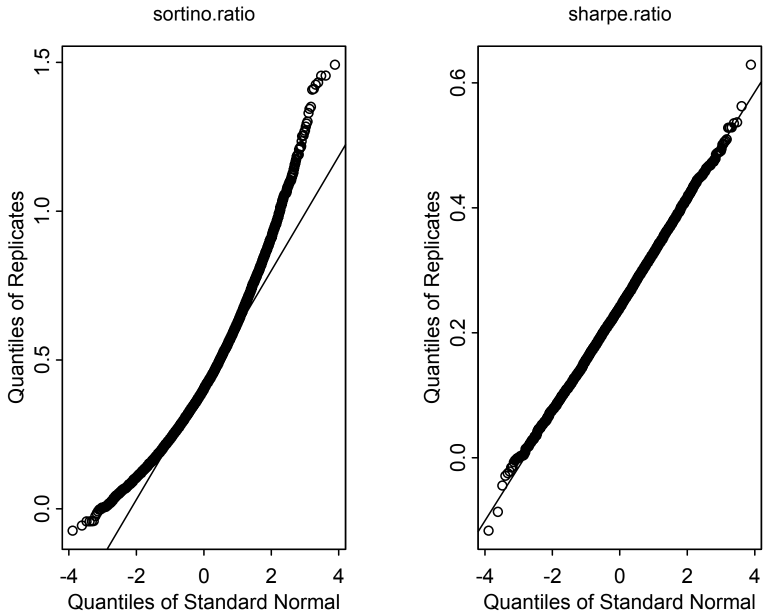

Figure 4.15 Q-Q Plots for Bootstrapped Sampling Distribution

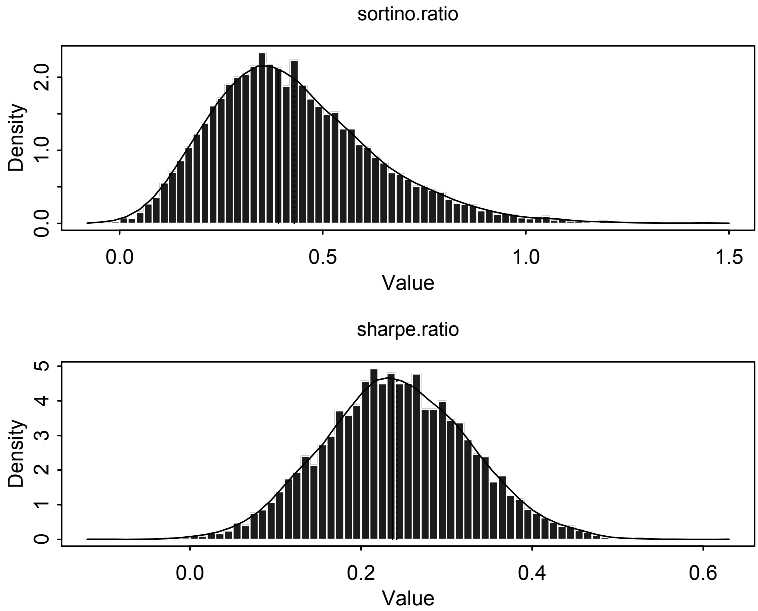

Figure 4.16 Bootstrapped Sampling Distribution

4.6.3 Increasing the Confidence Interval Coverage Probability Accuracy with the Double Bootstrap

So far we have relied on the 95% interval from a simple bootstrap procedure. However, the bootstrap is an approximate method, and it suffers to a greater or lesser extent from finite sample bias. Consequently, our 95% interval covers the true ratio with a probability that is at least somewhat different than .95. One way to increase the accuracy of the coverage probability of our confidence interval for the ratios is to use the **double bootstrap**, which can be thought of as "bootstrapping the bootstrap." It is known that under reasonable conditions the double bootstrap reduces the bias in the coverage probability.[13] The double bootstrap (Code 4.5) involves the following calculation.[14]

1. Perform the simple bootstrap as described above. Save all $b = 1,...,B$ resampled data sets as well as the resampled ratios $\hat{\varsigma}_b^*$. This is called *first stage-resampling*.
2. For each of the B resampled data sets, start a second round of $z = 1,\cdots,Z$ resamples, leading to a total of $B \cdot Z$ resamples denoted as $\hat{\varsigma}_{bz}^{**}$. For each $\hat{\varsigma}_b^*$, there exists a new set of Z resampled ratios $\hat{\varsigma}_{b1}^{**},\cdots,\hat{\varsigma}_{bZ}^{**}$. These are the *second-stage resamples*.
3. For each $\hat{\varsigma}_b^*$, calculate the percentage of second-stage resamples $\hat{\varsigma}_{b1}^{**},\cdots,\hat{\varsigma}_{bZ}^{**}$ that fall below the original sample estimate of the risk-return ratio $\hat{\varsigma}$, namely, calculate $u_b = \frac{1}{Z}\sum_{z=1}^{Z} I\left(\hat{\varsigma}_{bz}^{**} < \hat{\varsigma}\right)$. We choose $B = 1000$ and $Z = 200$.

```
double.bs <- function(data, statistic, B, Z)
{
  call(statistic)
  outer.sample <- matrix(
     sample(data, size=length(data)*B, replace=T),
     nrow=B, ncol=length(data))
  outer.bs <- apply(outer.sample, 1, statistic)
  inner.bs <- matrix(0, nrow=B, ncol=Z)
  prob <- matrix(0, ncol=1, nrow=B)
  estimate <- statistic(data)
  for(i in 1:B) {
     inner.bs[i,] <- bootstrap(outer.sample[i,],
        statistic, Z, trace=F)$replicates
     cat("run #", i)
```

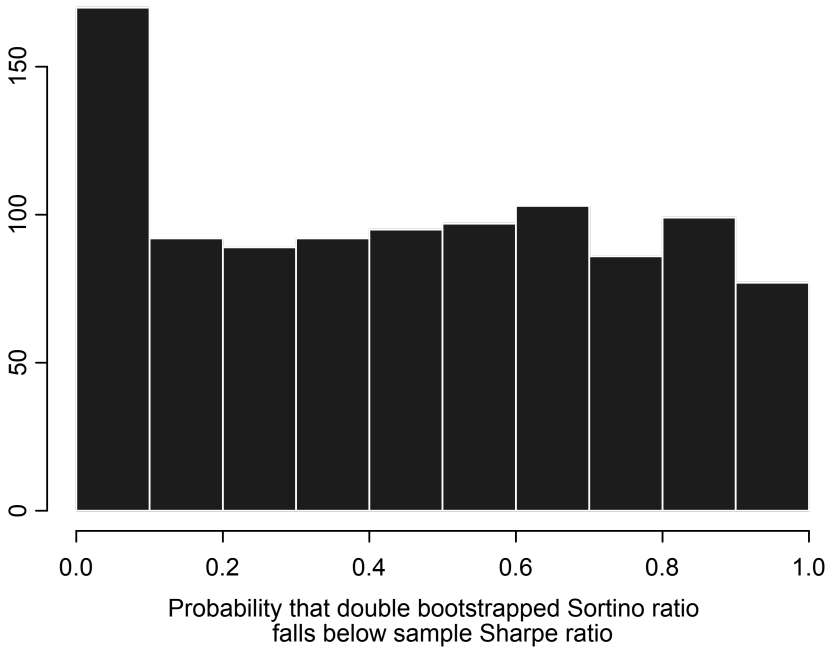

Figure 4.17 Sample Histogram of $u_b = \frac{1}{Z}\sum_{z=1}^{Z} I\left(\hat{\varsigma}_{bz}^{**} < \hat{\varsigma}\right)$

```
  }
  for(i in 1:B) {
     prob[i] <- sum(inner.bs[i,]<estimate)/Z
  }
  prob
}
```

Code 4.5 Double-Bootstrapping Code

Under ideal conditions, u_b follows a uniform distribution. Figure 4.17 shows that this assumption is clearly violated for the double-bootstrapped Sortino ratios $\hat{\varsigma}_{b1}^{**},\cdots,\hat{\varsigma}_{bZ}^{**}$.[15] Finally, we calculate the 2.5% and 97.5% percentiles of u_b and use these values to adjust the first-stage resample confidence band to $CI\left(\hat{\varsigma}_{u_{2.5\%}}^{*},\hat{\varsigma}_{u_{97.5\%}}^{*}\right)$. Our resulting double-bootstrap confidence interval $CI\left(\hat{\varsigma}_{7.9\%}^{*} = 0.18,\, \hat{\varsigma}_{96.\%}^{*} = 0.96\right)$ is moved to the right, with a higher lower bound of 0.18 instead of 0.11 (representing the 8% quantile rather than the 2.5% quantile), and a higher upper bound of .96 instead of .92.

Renewed interest in the significance of risk-return ratios has been focused on closed-form solutions for the well-known Sharpe ratio, and there is increasing interest in downside risk measures such as that in the Sortino ratio. This section provided a nonparametric methodology to evaluate the properties of the Sharpe and Sortino ratios' sampling distributions, as well as a method to compute confidence intervals without having to rely on asymptotic approximations. We have seen that while the distribution of the Sharpe ratio is well-approximated by a normal distribution, the Sortino ratio has a quite non-normal distribution, and the double-bootstrap methodology leads to a significantly refined confidence interval.

Exercises

1. This exercise points out an important linear regression model formulation of the Markowitz portfolio optimization without a long-only constraint, a context in which one can obtain standard errors of portfolio weights without resampling. Suppose we have n time series of excess returns (total return minus cash rate) with m observations each. We can combine these excess returns in a matrix $\mathbf{X}$ (each column contains one return series). Regressing these excess returns against a constant $\mathbf{1}_{mx1} = \mathbf{X}_{mxn}\mathbf{w}_{nx1} + \mathbf{u}_{mx1}$ yields $\mathbf{w} = \left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{X}^T\mathbf{1}$. These weights correspond to a portfolio that can be interpreted as the closest to a portfolio with zero risk (a vector of ones shows no volatility) and unit return. This would be an arbitrage opportunity. Rescaling the optimal weight vector (so that all weights sum to one) will yield $\mathbf{w}^*_{Sharpe} = \left(\bar{\mathbf{\Omega}}_0^{-1}\bar{\boldsymbol{\mu}}_0\right) \big/ \left(\mathbf{I}^T\bar{\mathbf{\Omega}}_0^{-1}\bar{\boldsymbol{\mu}}_0\right)$, the maximum Sharpe ratio portfolio. This framework can also be used to test restrictions on individual regression coefficients (estimated portfolio weights), as well as restrictions on groups of assets, and test whether they are significantly different from zero.[16]
   (a) Generate a hypothetical data set and use the linear regression command `lm()` in S-PLUS to calculate optimal portfolios.
   (b) Test for the significance of individual weights using alternative correlations and sample length.
   (c) Repeat (a), but add the constraints to the regression. Implement individual constraints, group constraints, and the full investment constraint.

2. Try to replicate Figures 4.4 to 4.9.

3. Make an equal-weighted portfolio of six to ten stocks of your choice from the CRSP returns data sets provided with this book, and apply the bootstrap and double bootstrap analysis of Section 4.6 to the Sharpe ratio and Sortino ratio for these data.

4. Repeat Exercise 3 for a new ratio obtained by modifying the Sortino ratio as follows: replace the denominator with the average of the losses below zero. How does the behavior of this ratio compare with that of the Sortino ratio?

5. Take the data from Michaud (1998, p.17, 19, given below in Table 4.1) and generate a graph similar to the graph in Figure 4.18.

Table 4.1 Data from Michaud (1998) for Exercise 5

$$\bar{\mathbf{\Omega}}_0 = \begin{array}{l} Canada \\ France \\ Germany \\ Japan \\ U.K. \\ U.S. \\ U.S.\ bonds \\ Euro\ bonds \end{array} \begin{pmatrix} 30.25 & 15.85 & 10.26 & 9.68 & 19.17 & 16.79 & 2.87 & 2.83 \\ 15.85 & 49.42 & 27.11 & 20.79 & 22.82 & 13.30 & 3.11 & 2.85 \\ 10.26 & 27.11 & 38.69 & 15.33 & 17.94 & 9.098 & 3.38 & 2.72 \\ 9.68 & 20.79 & 15.33 & 49.56 & 16.92 & 6.66 & 1.98 & 1.76 \\ 19.17 & 22.82 & 17.94 & 16.92 & 36.12 & 14.47 & 3.02 & 2.72 \\ 16.79 & 13.30 & 9.098 & 6.66 & 14.47 & 18.49 & 3.11 & 2.82 \\ 2.87 & 3.11 & 3.38 & 1.98 & 3.02 & 3.11 & 4.04 & 2.88 \\ 2.83 & 2.85 & 2.72 & 1.76 & 2.72 & 2.82 & 2.88 & 2.43 \end{pmatrix},$$

$$\bar{\boldsymbol{\mu}}_0 = \begin{array}{l} Canada \\ France \\ Germany \\ Japan \\ U.K. \\ U.S. \\ U.S.\ bonds \\ Euro\ bonds \end{array} \begin{pmatrix} 0.39 \\ 0.88 \\ 0.53 \\ 0.88 \\ 0.79 \\ 0.71 \\ 0.25 \\ 0.27 \end{pmatrix}$$

6. Redo the calculation in Section 4.6 with simulated data. What do you observe?

7. Select eight mid-cap stocks from `midcap.ts`, and compute the following resampled efficient frontiers: (a) resampling with the basic Michaud efficient frontier resampling described in Exercise 5; (b) resampling with a proper parametric bootstrap (i.e., evaluate each resampled portfolio mean and standard deviation by using the sample mean and covariance that generated the portfolio weights for that resampling, not the original sample mean and covariance as proposed by Jorion (1992) and Michaud (1998)); (c) the nonparametric bootstrap as described in Section 6.9.4, with simplified versions of the code provided in that section. What do you conclude about your results in (a) versus (b)? What about (b) versus (c)?

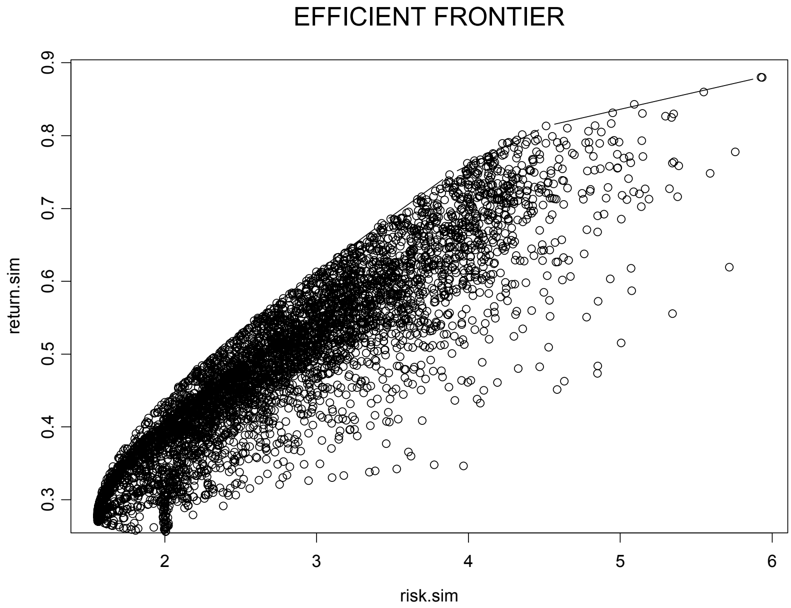

Figure 4.18 Efficient Frontier and Resampled Portfolio

Endnotes

[1] Robert Michaud patented the use of the average resampled portfolio. Readers are referred to U. S. patent # 6003018 or to Michaud (1998). However, the basic idea of portfolio resampling was introduced into the finance literature by Jorion (1992).
[2] See Efron and Tibshirani (1998) for a further discussion on this question.
[3] The assumption of a known fixed risk aversion coefficient is not always realistic, and if instead we use the weight vector $\mathbf{w}_i = \dfrac{\mathbf{\Omega}^{-1}\hat{\boldsymbol{\mu}}_i}{\mathbf{1}'\mathbf{\Omega}^{-1}\hat{\boldsymbol{\mu}}_i}$ with estimated risk aversion for the maximum Sharpe ratio for the i-th resample, we do not get this result.
[4] The central motivation of bootstrap resampling as introduced by statisticians is to estimate the distribution, or aspects of the distribution of an estimate such as the mean, standard deviation, or confidence intervals, of complicated statistics for which the standard sampling distribution theory does not apply.
[5] See Chapter 7 on Bayesian methods.
[6] The idea of this test statistic is that it is obviously not enough to look at weight differences only. Small weight differences for highly correlated assets might be of higher significance than large weight differences for negatively correlated assets.
[7] The Michaud approach referenced in Endnote 1 uses the rank-based approach.
[8] In spite of this apparent limitation, bootstrap resampling methods are able to do a quite decent job of estimating the distribution (or a summary such as standard error) of a statistic for which one does not have a decent sampling-distribution approximation; see, for example, Efron and Tibshirani (1998).
[9] See Lo (2002) and Memmel (2003).
[10] We assume here that returns are independently drawn from a single distribution. This is unlikely to be true for hedge fund data, as they exhibit serial correlation. One way to deal with this would be to fit an autoregressive model to the data and use this parametric specification of the return-generating process for resampling.
[11] See Sharpe (1994) for a review.
[12] See Sortino and Price (1994).
[13] See Section 3.9 and related material in Davison and Hinkley (1999) for details.
[14] See Nankervis (2002).
[15] Formal tests such as the Kolmogorov-Smirnov test as well as the χ^2 adjustment test, provide p-values close to 0%. Hence the null hypothesis that Figure 4.17 comes from a uniform distribution can be safely rejected.
[16] The regression framework puts a central problem of portfolio construction into a different, well-known perspective. Highly correlated asset returns mean highly correlated regressors with the obvious consequences arising from multicollinearity: high standard deviations on portfolio weights (regression coefficients) and identification problems (difficulty of distinguishing between two similar assets). Simply downtesting and excluding insignificant assets will result in an outcome that is highly dependent on the order of exclusion, with no guidance where to start. This is a familiar problem for both the asset allocator and the econometrician.