

Markowitz portfolio optimization using MOSEK.

MOSEK Technical report: TR-2009-2

Erling D. Andersen and Joachim Dahl*

Revised on February 9th, 2011.

Abstract

In this tutorial paper we introduce different approaches to Markowitz portfolio optimization, and we show how to solve such problems in MATLAB using the MOSEK optimization toolbox for MATLAB. We first consider conic formulations of the basic portfolio selection problem, and we then discuss more advanced models for transaction costs, as well as non-convex cardinality constraints.

1 Markowitz portfolio selection

We start by reviewing basic Markowitz portfolio optimization, while introducing the necessary notation used in the remaining tutorial. In Markowitz portfolio selection we optimize a portfolio of stock assets based on a simple probabilistic model of the stock market. Traditionally, we assume that the return of n different assets over time can be modeled as a multivariate random variable $r \in \mathbf{R}^n$ with known mean

$$\mathbf{E} r = \mu_r$$

and covariance matrix

$$\mathbf{E}(r - \mu_r)(r - \mu_r)^T = \Sigma_r,$$

where the mean is a vector, and Σ_r is the covariance matrix (i.e., $\mu_r \in \mathbf{R}^n$ and $\Sigma \in \mathbf{R}^{n \times n}$). We have an initial holding of w_j^0 dollars of asset j , $j = 1, \dots, n$ and we invest x_j dollars in asset j , i.e., after the investment period our portfolio is $w^0 + x$. The return of our investment is also a random variable,

$$y = r^T(w^0 + x)$$

with mean value (or *expected return*),

$$\mathbf{E} y = \mu_r^T(w^0 + x) \tag{1}$$

and variance (or *risk*),

$$(y - \mathbf{E} y)^2 = (w^0 + x)^T \Sigma_r (w^0 + x). \tag{2}$$

In Markowitz portfolio selection, we seek to optimize a trade-off between expected return and risk.

*MOSEK ApS, Fruebjergvej 3, Box 16, 2100 Copenhagen, Denmark. Email: support@mosek.com

1.1 Estimating the mean and variance

In practice, the mean return μ_r and covariance Σ_r must be estimated based on empirical data, which is often a non-trivial task; in this tutorial we estimate those entities using simple *sample estimates*. More details on estimating the parameters from historical data are given in Appendix C.

We assume that for all n assets, we have recordings of N stock returns at different times, and we organize those measurements in a matrix $X \in \mathbf{R}^{N \times n}$, where column j represents the return of asset j over time. This is similar to how such data would be represented in a spreadsheet with n columns and N rows. In MATLAB we can easily estimate the sample mean \bar{r} and covariance $\bar{\Sigma}_r$ as shown below:

```
1 % Estimate sample mean and variance of X
  mu_r = mean(X, 1);
3 Sigma_r = cov(X);
```

A recurring theme in this tutorial is *conic quadratic* formulations (see Appendix A), which are well-suited for the MOSEK optimizer. In the following conic formulations of Markowitz portfolio selection, we need a factorization of the covariance Σ_r . The covariance matrix can be factored as

$$\bar{\Sigma}_r = G^T G,$$

where the G matrix depends on the choice of factorization. For any such factorization, we can then express the risk as a Euclidean norm, i.e.,

$$(w^0 + x)^T \bar{\Sigma}_r (w^0 + x) = (w^0 + x)^T G^T G (w^0 + x) = \|G(w^0 + x)\|^2,$$

which is used in the conic formulations in the following sections. The sample estimate is defined as a factorization (or a *Grammian matrix*)

$$\bar{\Sigma}_r \triangleq \bar{X}^T \bar{X},$$

where

$$\bar{X} \triangleq \frac{1}{\sqrt{N-1}}(X - e\bar{r}^T)$$

denotes a scaled zero-mean version of the data-matrix, and e is a vector of all ones. We can therefore express the risk directly in terms of the (normalized) data-matrix \bar{X} as

$$(w^0 + x)^T \bar{\Sigma}_r (w^0 + x) = (w^0 + x)^T \bar{X}^T \bar{X} (w^0 + x) = \|\bar{X}(w^0 + x)\|^2. \quad (3)$$

The factored risk expression in (3) is the one most frequently used in Markowitz portfolio optimization, and makes no assumptions about the dimensions or the rank of X - it can be employed independent of we have more observations than assets (i.e., whether X has more rows than columns). Moreover, it is not required that X has full rank. In MATLAB the normalized data-matrix can be computed as shown below:

```
1 % Normalize data-matrix for conic formulations of risk
  [N, n] = size(X);
3 mu_r = mean(X, 1);
  Xbar = 1/sqrt(N-1)*(X-ones(N,1)*mu_r');
```

Alternatively, we can use a QR factorization on the normalized data-matrix

$$QR = \bar{X}$$

where $Q \in \mathbf{R}^{N \times N}$ is an orthogonal matrix (i.e., $QQ^T = Q^TQ = I$) and $R \in \mathbf{R}^{N \times n}$ is an upper-triangular matrix. Since Q is orthogonal, we have that

$$(w^0 + x)^T \bar{\Sigma}_r (w^0 + x) = (w^0 + x)^T R^T Q^T Q R (w^0 + x) = \|R(w^0 + x)\|^2.$$

In MATLAB the QR factorization can be computed as:

```

2 % Compute a factorization of the normalized data-matrix
[Q, R] = qr(Xbar); G = R;
```

A related approach to the QR factorization is the *singular-value decomposition* (SVD) of \bar{X} ,

$$\bar{X} = USV^T,$$

where $U \in \mathbf{R}^{N \times N}$ is an orthogonal matrix, $S \in \mathbf{R}^{N \times n}$ is a diagonal matrix with non-negative *singular values* appearing in decreasing order on the diagonal

$$S_{11} \geq S_{22} \geq \dots \geq S_{pp} \geq 0,$$

where p is the rank of \bar{X} , and $V \in \mathbf{R}^{n \times n}$ is another orthogonal matrix. As for the QR factorization, we get a factored expression for the risk,

$$(w^0 + x)^T \bar{\Sigma}_r (w^0 + x) = (w^0 + x)^T V S^T U^T U S V^T (w^0 + x) = \|S V^T (w^0 + x)\|^2.$$

In MATLAB the SVD can be computed as:

```

2 % Compute a singular value decomposition of the normalized data-matrix
[U, S, V] = svd(Xbar); G = S*V';
```

The SVD is often used for creating *low-rank approximations*, i.e., by truncating the smallest singular values to zero, we obtain a low-rank approximation \hat{X} of \bar{X} . This is often useful for robust modeling and for removing noise from the measurements, but is beyond the scope of this tutorial.

We conclude this section by mentioning the simpler *Cholesky* factorization, which can be used to factor symmetric positive definite matrices, e.g., if $\bar{\Sigma}_r$ is positive definite, we can factor it as

$$\bar{\Sigma}_r = R^T R,$$

where R is an upper-triangular matrix. In MATLAB we compute the Cholesky factorization as:

```

2 % Compute a Cholesky factorization of the covariance matrix
Sbar = Xbar'*Xbar;
R = chol(Sbar);
```

In general, it is not recommended to use the Cholesky factorization approach for factoring Σ_r instead of directly using the (normalized) data matrix, or a QR factorization of the data-matrix. When we have fewer observations than assets, the data matrix has smaller dimensions than the covariance matrix, so from a computational point of view it is more attractive. Furthermore, in that case, the covariance matrix is rank-deficient, and as such the Cholesky factorization is not directly applicable. On the other hand, when we have more observations than assets, the covariance matrix will generally have full rank and its Cholesky factor has smaller dimensions than the data matrix.

The economy size QR factorization unifies the two cases; it yields the smallest factor, and can be used regardless of the rank of the data-matrix and covariance matrix. A more subtle benefit of the QR factor compared to the Cholesky factor is that it avoids the squaring operation for forming the covariance matrix, and subsequently it is better conditioned. However, in Section 1.6 the Cholesky factorization will prove very useful.

1.2 Minimizing the risk

We now turn to the most traditional Markowitz portfolio selection problem, where we minimize the risk (in this case by minimizing the square root of the variance) given a fixed expected return,

$$\begin{aligned} & \text{minimize} && \|R(w^0 + x)\| \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t, \\ & && e^T x = 0. \end{aligned} \tag{4}$$

The additional constraint $e^T x = 0$ is a *budget* or *self-financing* constraint stating that the total amount invested in the portfolio remains unchanged after the transaction.

To rewrite (4) as a standard conic quadratic problem, we first use an intuitive reformulation; the problem

$$\text{minimize } g(x)$$

can be formulated equivalently in *epigraph form* using one additional variable f as,

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && g(x) \leq f. \end{aligned}$$

We can therefore rewrite (4) as,

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t, \\ & && \|R(w^0 + x)\| \leq f, \\ & && e^T x = 0. \end{aligned}$$

The problem can finally be formulated as a standard conic quadratic optimization problem suited for MOSEK by including an additional equality constraint as

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t, \\ & && R(w^0 + x) = g, \\ & && f \geq \|g\|, \\ & && e^T x = 0 \end{aligned} \tag{5}$$

in the variables (x, f, g) .


Code listing 1 shows a MATLAB example program for minimum risk selection for the S&P500 data set described in Appendix C for 500 stocks with 800 observations (i.e., $X \in \mathbf{R}^{800 \times 500}$). We initially have 1 dollar distributed over the portfolio, i.e., $w^0 = (1/n)e$, and we demand an expected daily return of $t = 1.0005$. If we run the program, we get an optimal portfolio illustrated in Fig. 1.2a. From the example we observe an investment strategy called *short-selling*. With short-selling, the investor borrows the asset and sells it, and at the end of the period the asset is bought back and returned. Mathematically, short-selling occurs whenever

$$w_j^0 + x_j \leq 0,$$


and if the asset has lost value during the transaction period, the investor makes a profit buying it back.

The potential loss associated with short-selling is infinite, and therefore we often wish to either eliminate or limit short-selling. Short-selling can be completely eliminated from the portfolio selection by adding a linear constraint

$$w^0 + x \geq 0,$$



(a) With short-selling. The associated risk is $f = 0.0068$.



(b) Without short-selling. The associated risk is $f = 0.0836$.

Figure 1: Optimal minimum-risk solutions.

or we can allow a limited short-selling of s_j for asset j ,

$$w_j^0 + x_j \geq -s_j.$$

If we modify the program in listing 1 to exclude short-selling (by changing line 18), we get a very different investment strategy illustrated in Fig. 1.2b.

Since we are restricting our feasible domain by adding the constraints $w_k^0 + x_j \geq 0$, we invariably increase the objective value, i.e., the risk. With this investment strategy, we must invest heavily in a few assets to achieve our expected return. This naturally leads to the question: how high an expected return can we get? With short-selling, we can achieve an arbitrarily high expected return. Without short-selling, the largest expected return occurs when we invest exclusively in a single asset - in this example in asset #481. If we change the demanded return to

$$t = \max(r) + 1;$$

and rerun the example, we get the following message

```
Interior-point solution
Problem status : PRIMAL_INFEASIBLE
Solution status : PRIMAL_INFEASIBLE_CER
```

which tells us that the constraints are conflicting (i.e., the feasible set is empty); such a large expected return is incompatible with the self-budget constraint.

Another pertinent issue is which factorization of the data matrix do we use. In the example we considered a time-horizon of 800 observations, i.e., $X \in \mathbf{R}^{800 \times 500}$ and $\Sigma = X^T X \in \mathbf{R}^{500 \times 500}$ has full rank, so the most efficient formulation (in terms of solution time) uses an economy-size QR factorization of the data-matrix, $X = QR$ with a triangular 500×500 factor R . If instead we directly use the data matrix X as our factor, the solution time for this particular example is around 6 times slower¹.

The situation is different when we have a short time-horizon with fewer observations than number of stocks. In that case, there is no efficiency loss in directly using the data matrix instead of the economy-size QR factorization (since then the dimensions of X and R are identical), and using the data-matrix is simpler and saves a single QR factorization. If we change the time-horizon in the example to 100 observations, then the solution time using either the QR factorization or the data matrix is around 7 times faster than with 800 observations. We summarize the execution times for the different setups in Table 1.

# observations	factor	solution time (secs)
800	QR	1.09
800	data matrix	6.47
100	QR	0.16
100	data matrix	0.16

Table 1: Solution times for different portfolio setups.

As a final experiment, we keep the number of observations fixed at $N = 800$, but we gradually change the number of stocks in the portfolio from 50 to 500 and record the solution. Fig. 2 shows the solution time for different problem sizes.

¹All examples are executed on a quad core Intel Xeon 2.5GHz processor with 16GB memory.




Figure 2: Solution time for different problem sizes. The dashed line shows the fitted curve $f(x) = \gamma x^{2.66}$, where γ depends on the computer (we observed $\gamma \approx 10^{-7}$).

1.3 Maximizing expected return

We can also maximize the expected return given constraints on the largest admissible risk,

$$\|R(w^0 + x)\| \leq \hat{f},$$

where \hat{f} is the acceptable level of risk. This leads to the problem

$$\begin{aligned} &\text{maximize} && \bar{r}^T(w^0 + x) \\ &\text{subject to} && e^T x = 0, \\ &&& \|R(w^0 + x)\| \leq \hat{f}, \end{aligned}$$

which we can rewrite as a standard conic quadratic problem

$$\begin{aligned} &\text{maximize} && \bar{r}^T(w^0 + x) \\ &\text{subject to} && e^T x = 0, \\ &&& R(w^0 + x) = g \\ &&& f \geq \|g\|, \\ &&& f = \hat{f} \end{aligned} \tag{6}$$

in the variables (x, f, g) .

1.4 Computing the efficient frontier

The two problem formulations (5) and (6) produce the same *optimal frontier*, i.e., by varying t and \hat{f} the two formulations give the same optimal trade-off curves between risk and expected return. If our main objective is to compute the efficient frontier, we can use an alternative

formulation that directly minimizes a compromise between standard deviation and return using a positive *penalization factor* λ ,

$$\begin{aligned} & \text{maximize} && \bar{r}^T(w^0 + x) - \lambda \|R(w^0 + x)\| \\ & \text{subject to} && e^T x = 0, \end{aligned}$$

which has the following conic quadratic formulation

$$\begin{aligned} & \text{maximize} && \bar{r}^T(w^0 + x) - \lambda f \\ & \text{subject to} && e^T x = 0, \\ & && R(w^0 + x) = g \\ & && f \geq \|g\|, \end{aligned} \tag{7}$$

in the variables (x, f, g) . Alternatively, we may wish to compute an efficient trade-off between risk and return (with the same unit, e.g., dollars),

$$\begin{aligned} & \text{maximize} && \bar{r}^T(w^0 + x) - (\lambda/2) \|R(w^0 + x)\|^2 \\ & \text{subject to} && e^T x = 0, \end{aligned}$$

which can be posed as a *rotated conic quadratic* problem (see Appendix A),

$$\begin{aligned} & \text{maximize} && \bar{r}^T(w^0 + x) - f \\ & \text{subject to} && e^T x = 0, \\ & && R(w^0 + x) = g \\ & && 2fh \geq \|g\|^2 \\ & && h = 1/\lambda, \end{aligned} \tag{8}$$

in the variables (x, f, g, h) . By varying $\lambda \in [0, \infty)$ in either (7) or (8) we trace the (same) efficient frontier.

1.5 Maximizing the Sharpe ratio

The Sharpe ratio defines an efficiency metric of a portfolio as the expected return per unit risk, i.e.,

$$S(x) = \frac{r^T(w^0 + x) - r_f e^T w^0}{\|R(w^0 + x)\|},$$

where r_f denotes the return of a risk-free asset, and x is assumed not to include a risk-free asset. Also, by assumption, we have that $\bar{r}^T(w^0 + x) > r_f e^T w^0$ (i.e., there exists a risk-associated portfolio with a greater yield than the risk-free asset), so maximizing the Sharpe ratio is equivalent to minimizing $1/S(x)$. In other words, we have the following problem

$$\begin{aligned} & \text{minimize} && \frac{\|R(w^0 + x)\|}{\bar{r}^T(w^0 + x) - r_f e^T w^0} \\ & \text{subject to} && e^T x = 0. \end{aligned}$$

We next introduce a variable transformation

$$y = \gamma(w^0 + x), \quad \gamma \geq 0.$$

Since $\gamma > 0$ can be chosen arbitrarily and $(\bar{r} - r_f e)^T(w^0 + x) > 0$, we have without loss of generality that

$$(\bar{r} - r_f e)^T y = 1.$$

Thus, we obtain the following problem for maximizing the Sharpe ratio,

$$\begin{aligned} & \text{minimize} && \|Ry\| \\ & \text{subject to} && e^T y = \gamma \cdot e^T w^0 \\ & && (\bar{r} - r_f e)^T y = 1 \\ & && \gamma \geq 0, \end{aligned}$$

or equivalently as a standard conic quadratic problem,

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && e^T y = \gamma \cdot e^T w^0 \\ & && Ry = g \\ & && f \geq \|g\| \\ & && (\bar{r} - r_f e)^T y = 1 \\ & && \gamma \geq 0, \end{aligned}$$

and we recover $x := y/\gamma - w^0$. Eliminating short-selling corresponds to adding the constraints $y \geq 0$ (in fact $y/\gamma \geq 0$). Listing 2 shows a MATLAB implementation for computing the largest Sharpe ratio without short-selling.

We next illustrate the largest Sharpe ratio with a numerical example. For the data set used in §1.2 (i.e., a portfolio of 500 assets and 800 observations) we compute the efficient frontier (using the algorithm in §1.2 for different values of t), and we also compute the largest Sharpe ratio, assuming a risk-free asset with a daily return of 1.0005.

The results are shown in Fig. 3. As evident from the figure, the maximum Sharpe ratio is the tangent to the optimal frontier, which passes through $(0, r_f)$.




Figure 3: Efficient frontier of a portfolio. The solution that achieves the maximum Sharpe ratio is denoted by ‘o’.

1.6 Structured covariance matrices

Sometimes the covariance matrix is assumed to have a specific structure, e.g., in a *factor model* it is often assumed to consist of a diagonal plus a low-rank term,

$$\Sigma_r = D + ABA^T \tag{9}$$

where $D \in \mathbf{R}^{n \times n}$ is a positive definite diagonal matrix, $B \in \mathbf{R}^{l \times l}$ ($l \ll n$) is a positive definite matrix, and $A \in \mathbf{R}^{n \times l}$ is a matrix that describes the impact of different factors on the covariance. We can think of (9) as a simple *parametric model*, with parameters D , A and B ,

and since the number of parameters in (9) is smaller than the number of parameters in $\tilde{\Sigma}_r$, we expect the parameters in the factor model to be estimated more accurately. Popular factor models are developed by commercial risk analysis companies, and the parameters are estimated using expert knowledge of the respective stock market.

We could use a Cholesky factorization from Section 1.1 to factor Σ_r directly, which would result in a triangular dense factor. Alternatively, we can rewrite the factor model as

$$\Sigma_r = \begin{bmatrix} I & A \end{bmatrix} \begin{bmatrix} D & \\ & B \end{bmatrix} \begin{bmatrix} I & A \end{bmatrix}^T,$$

and factor the smaller matrix $B = R^T R$ using a Cholesky factorization. If we define

$$G \triangleq \begin{bmatrix} D^{1/2} & \\ & R \end{bmatrix} \begin{bmatrix} I & A \end{bmatrix}^T = \begin{bmatrix} D^{1/2} \\ RA^T \end{bmatrix}$$

where $D^{1/2}$ denotes a positive diagonal matrix with $D_{ii}^{1/2} = \sqrt{D_{ii}}$, we get an alternative factorization $\Sigma_r = G^T G$. Thus, exploiting the structure of the factor model leads to a problem

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \tilde{r}^T(w^0 + x) = t, \\ & && G(w^0 + x) = \tilde{g}, \\ & && f \geq \|\tilde{g}\|, \\ & && e^T x = 0 \end{aligned}$$

in the variables (x, f, \tilde{g}) . The dimensions of this formulation are larger than for the original problem (5), but the constraints are much sparser (and therefore require less storage space), which generally is preferable.

1.7 Choosing the right formulation

In this section we summarize important points made in previous sections about speeding up the solution time.

- In general, one should use the formulation with the smallest memory requirements.
- Exploiting the low-rank structure of a factor model gives a significant speedup in solution time, even though we get additional variables. In the quadratic programming formulation, we can think of the improvement as replacing a completely dense Hessian matrix of the objective with a larger diagonal Hessian plus a number of relatively sparse constraints.
- If we need to estimate the covariance matrix using fewer observations than assets (i.e., we have a relatively short time-horizon), then we should directly use the data-matrix as our factor, or alternative we can use an economy size QR factoriation.
- If we need to estimate the covariance matrix, and we have more observations than assets (i.e., we have a relatively long time-horizon), then we should use the economy size QR factorization of the data matrix, or alternative we can use a Cholesky factorization of the estimated covariance matrix, but we should not directly use the data matrix as our factor.

2 Transaction costs with market impact

The previous problem formulations in Section 1 did not include transaction costs, i.e., they assumed no cost associated with a transaction, and that trading does not affect the market

price. A more realistic model of the transaction cost (of a single asset) is given in [1], and has the form

$$\text{commission} + \frac{\text{bid}}{\text{ask}} - \text{spread} + \theta \sqrt{\frac{\text{trade volume}}{\text{daily volume}}}.$$

The last term is the *market impact cost*, and captures that the price of an asset increases or decreases if you buy or sell large quantities, respectively. It can be approximated as

$$\theta_j \sqrt{\frac{(\text{trade volume})_j}{(\text{daily volume})_j}} \approx m_j \sqrt{|x_j|}$$

for the j th asset, where θ_j and m_j are parameters that must be estimated. We can include the market impact term into the minimum-risk portfolio formulation as follows

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t, \\ & && \|R(w^0 + x)\| \leq f, \\ & && e^T x + e^T y = 0, \\ & && |x_j|(m_j|x_j|^{1/2}) \leq y_j, \quad j = 1, \dots, n, \end{aligned}$$

where the market impact term influences the modified self-budget constraint $e^T x + e^T y = 0$ to make large transaction more costly. If we introduce the variable change $y_j = m_j \bar{y}_j$, we get the formulation

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t, \\ & && \|R(w^0 + x)\| \leq f, \\ & && e^T x + m^T \bar{y} = 0, \\ & && |x_j|^{3/2} \leq \bar{y}_j, \quad j = 1, \dots, n. \end{aligned} \tag{10}$$

The formulation (10) is a convex problem, but more interestingly, it can also be written as a standard conic quadratic problem; the derivation is technically more involved and deferred to Appendix B. Using the result from Appendix B, we get the following standard conic quadratic formulation of the minimum risk selection with market impact correction,

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t \\ & && e^T x + m^T \bar{y} = 0 \\ & && R(w^0 + x) = g \\ & && f \geq \|g\| \\ & && z_j^2 \leq 2s_j \bar{y}_j \quad j = 1, \dots, n \\ & && w_j^2 \leq 2v_j r_j \quad j = 1, \dots, n \\ & && z = v \\ & && s = w \\ & && r = \frac{1}{8}e \\ & && x \leq z \\ & && -x \leq z \\ & && s, \bar{y}, v, r \geq 0, \end{aligned} \tag{11}$$

in the variables $(x, f, g, s, \bar{y}, v, r, z)$. The conic formulation contains several additional variables and constraints, but the constraints are all extremely sparse, so they only have a very limited impact on memory requirements and execution time for the MOSEK solver.

3 Cardinality constrained portfolio selection

The transaction cost can be modeled in different ways, e.g., we can assume a linear cost

$$T(x_j) = \begin{cases} c^+ x_j, & x_j \geq 0 \\ -c^- x_j, & x_j < 0, \end{cases}$$

or more realistically, we can assume a constant transaction cost plus a linear cost,

$$T(x_j) = \begin{cases} b + c^+ x_j, & x_j > 0 \\ 0, & x_j = 0 \\ b - c^- x_j, & x_j < 0. \end{cases}$$

The latter model can be included in the portfolio selection model as

$$\begin{aligned} & \text{minimize} && f \\ & \text{subject to} && \bar{r}^T(w^0 + x) = t \\ & && R(w^0 + x) = g \\ & && f \geq \|g\| \\ & && e^T x + e^T y = 0 \\ & && bz_j + c^+ x_j \leq y_j \\ & && bz_j - c^- x_j \leq y_j \\ & && x_j \geq l_j z_j \\ & && x_j \leq u_j z_j \\ & && z_j \in \{0, 1\} \end{aligned} \tag{12}$$

in the variables (x, y, f, g, z) . In the formulation (12), the variable y_j is the transaction cost for x_j and the binary variable z_j controls the constant transaction cost. The constants l_j and u_j give lower and upper bounds on the allowed range of x_j , respectively. This problem is non-convex because of the integer constraints $z_j \in \{0, 1\}$, and is therefore much more difficult to solve than the problems studied in the previous sections. Tight bounds l_j and u_j generally help to reduce computation-time, however.

A Conic quadratic optimization

In this section, we review conic quadratic optimization. Conic quadratic optimization is an established field of convex optimization with close similarities to traditional linear programming. A conic quadratic optimization problem can be formulated as

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Ax = b \\ & && x \in \mathcal{C}_q \end{aligned} \tag{13}$$

in the variable $x \in \mathbf{R}^n$ with a problem specified by $c \in \mathbf{R}^n$, $A \in \mathbf{R}^{m \times n}$ and $b \in \mathbf{R}^m$. The cone \mathcal{C}_q is the so-called *quadratic cone* or *Lorenz cone*, and $x \in \mathcal{C}_q$ means that

$$x_1 \geq \sqrt{x_2^2 + \cdots + x_n^2}.$$

Note, that this formulation is very similar to a standard linear program

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Ax = b \\ & && x \in \mathcal{C}_l, \end{aligned} \tag{14}$$

where the *linear cone* \mathcal{C}_l is simply the non-negative orthant, i.e., $\mathcal{C}_l = \{x \mid x \geq 0\}$.

MOSEK is an industry leader in software for conic optimization, and uses a more general formulation for the conic quadratic problem

$$\begin{aligned} & \text{minimize} && c^T x + c^f \\ & \text{subject to} && l^c \leq Ax \leq u^c \\ & && l^x \leq x \leq u^x \\ & && x^t \in \mathcal{C}^t \end{aligned} \tag{15}$$

where the variables x^t form a partitioning of \mathbf{R}^n , i.e., each element of x belongs to *exactly one* subset x^t . For example, it could be that

$$x^1 = \begin{bmatrix} x_1 \\ x_4 \\ x_7 \end{bmatrix}, \quad x^2 = \begin{bmatrix} x_6 \\ x_5 \\ x_3 \\ x_2 \end{bmatrix}.$$

Each cone \mathcal{C}^t is defined as either,

- the set of *free variables*

$$\mathcal{C}_f^t = \{x \in \mathbf{R}^{n^t}\},$$

- the *quadratic cone*

$$\mathcal{C}_q^t = \left\{ x \in \mathbf{R}^{n^t} : x_1 \geq \sqrt{\sum_{j=2}^{n^t} x_j^2} \right\},$$

- the *rotated quadratic cone*

$$\mathcal{C}_r^t = \left\{ x \in \mathbf{R}^{n^t} : 2x_1x_2 \geq \sum_{j=3}^{n^t} x_j^2, x_1, x_2 \geq 0 \right\}.$$

Note, that the formulation (15) covers the standard conic quadratic formulation (13) by choosing $l^c = u^c = b$, $l^x = -\infty$, $u^x = \infty$ and a single quadratic cone $x \in \mathcal{C}_q$. Similarly, it also covers the standard linear programming formulation (14) by choosing $l^c = u^c = b$, $l^x = 0$, $u^x = \infty$ and a single cone $x \in \mathcal{C}_f$. More generally, the conic quadratic formulation (15) covers a wide range of different problem formulations and applications, see [3, 4] for a general treatment.

B Conic quadratic modeling of $|x|^{3/2}$

In this appendix we show how to model

$$|x|^{3/2} \leq t$$

as a conic quadratic constraint. We initially represent it as

$$-z \leq x \leq z, \quad \frac{z^2}{\sqrt{z}} \leq t, \quad z \geq 0$$

using an additional variable z . The only non-linear constraint is $z^2/\sqrt{z} \leq t$, which can be represented equivalently as

$$\begin{aligned} z^2 &\leq 2st \\ w^2 &\leq 2vr \\ z &= v \\ s &= w \\ r &= \frac{1}{8} \\ s, t, v, r &\geq 0, \end{aligned} \tag{16}$$

in the variables (z, w, s, t) . The set (16) is the intersection of a set of linear (in)equalities and rotated quadratic cones (see Appendix A), and thus fits the framework of conic quadratic optimization.

To see that the formulations are equivalent, we observe from (16) that

$$s^2 = w^2 \leq 2vr = \frac{v}{4} = \frac{z}{4}$$

and

$$z^2 \leq 2st \leq 2\sqrt{\frac{z}{4}}t,$$

which shows the desired result

$$z^2/\sqrt{z} \leq t.$$

Thus, a conic quadratic formulation of the inequality $|x|^{3/2} \leq t$ is

$$\begin{aligned} x &\leq z \\ -x &\leq z \\ z^2 &\leq 2st \\ w^2 &\leq 2vr \\ z &= v \\ s &= w \\ r &= \frac{1}{8} \\ s, t, v, r &\geq 0. \end{aligned} \tag{17}$$

More generally, for positive integers p and q the following convex power functions

$$|x|^{p/q} \leq t, \quad p/q \geq 1,$$

and

$$|x|^{-p/q} \leq t,$$

can be represented using conic quadratic formulations, see [3, 2] for details.

C Parameter estimation from historical data

In practice, the parameters of the stochastic portfolio model are not known and must be estimated from historical data. In this tutorial we consider daily prices for the 500 stocks in the S&P500 index [6].

For these 500 stocks, we retrieved historical prices from the Google Finance service [5]. The dataset contains daily observations of the opening, closing, highest, and lowest trading prices, respectively, as well as the traded volume. The number of observations for each stock varies, since stocks might enter or leave the index at different times, and furthermore, there are missing observations in the dataset.

From the data, we need to estimate the vector of average returns \bar{r} as well as the covariance matrix for the returns. For each stock asset j , we first estimate the k th daily absolute return as

$$X_{j,k} = \frac{(\text{closing price})_{j,k+1}}{(\text{closing price})_{j,k}}.$$

This gives us an $N \times 500$ data matrix with daily absolute returns, where N is the width of our observation window (or time-horizon). Estimating the average returns from the data matrix is by no means trivial. For example, consider two consecutive returns of 2 and 1/2. A reasonable estimate of the average return is the geometric mean of 1, but the arithmetic mean of 5/4 is less appropriate. Furthermore, for realistic data spanning a longer time-horizon, we also need to consider other issues such as *discounting*, etc., and possibly use different weights for the observations.

Another issue is how to handle missing data in the observations. We choose to simply interpolate (and extrapolate) missing data, which is easily performed in MATLAB using the `interp1` function.

```
1   t = 1:length(X);
2   if any(isnan(X))
3       mask = ~isnan(X);
4       X(~mask) = interp1(t(mask),X(mask),t(~mask),'nearest','extrap');
5   end
```

To estimate the average return, we merely use the arithmetic mean of each column in X as our estimate, i.e.,

$$\bar{r} = (1/N)X^T e.$$

Similarly, we estimate the covariance of returns as the traditional unbiased estimator

$$\bar{\Sigma}_r = \frac{1}{N-1}(X - e\bar{r}^T)^T(X - e\bar{r}^T).$$

As an illustration, Fig. 4 shows the closing price for AAPL for the past 800 observations (or roughly 3.5 years). From the data we estimate an average daily return of 1.0015.




Figure 4: Closing prices for AAPL from 2007-12-07 to 2011-01-20.

D MATLAB programs

```

1 % Minimum risk Markowitz portfolio selection example
2 %
3 % minimize f
4 % subject to r'*(w0 + x) = t
5 %             G*(w0 + x) = g, ||g||_2 < f
6 %             sum(x) = 0.
7 %             and optionally w0 + x >= 0
8 % or equivalently
9 %
10 % minimize f
11 % subject to [ r'  0  0 ] [ x ] = [ t-r'*w0 ]
12 %            [ G  0 -I ]*[ f ] = [ -G*w0 ]
13 %            [ e'  0  0 ] [ g ] = [ 0 ]
14 %            ||g||_2 <= f
15 %            and optionally w0 + x >= 0
16 clear
17
18 SHORTSELLING = 0;
19 load sp500
20 X = R(1:800, :);
21
22 [N, n] = size(X);
23 r = mean(X)';
24 Xbar = 1/sqrt(N-1)*(X-ones(N,1)*r');
25
26 % Choose factorization
27 %G = Xbar;
28 [Q, G] = qr(Xbar, 0);
29 %[U, S, V] = svd(Xbar); G = S*V';
30
31 p = size(G,1)
32 w0 = ones(n,1)/n; % initial holdings
33 t = 1.005; % expected return
34
35 prob.a = sparse([ r', zeros(1, 1+p);
36                 G, zeros(p,1), -eye(p);
37                 ones(1,n), zeros(1,1+p)]);
38 prob.c = [ zeros(n,1); 1; zeros(p,1) ];
39 prob.cones{1}.type = 'MSK_CTQUAD';
40 prob.cones{1}.sub = [n+1:n+1+p];
41 b = [ t - r'*w0; -G*w0; 0 ];
42 prob.blc = b;
43 prob.buc = b;
44
45 if SHORTSELLING == 0
46     prob.blx = [-w0; -inf*ones(1+p,1)];
47 end
48 [err, res]=mosekopt('minimize', prob);
49 x = res.sol.itr.xx(1:n);
50 f = res.sol.itr.xx(n+1);

```

Listing 1: Minimum risk portfolio selection.

```

function [x, risk, ret] = sharpe(G, r, w0, rf)
2
% Computes the largest Sharpe ratio (assuming no short-selling),
4 %
% minimize      g
6 % subject to  (r-e*rf)'*y = 1
%              G*y = f,  ||f||_2 < g
8 %              sum(y) = t*sum(w0)
%              t >= 0
10 %             y >= 0
%
12 % or equivalently
%
14 % minimize      g
% subject to  [ (r-e*rf)'  0  0  0 ] [ y ] [ 1 ]
16 %            [ G      -I  0  0 ] * [ f ] = [ 0 ]
%            [ e'      0  0 -t*sum(w0) ] [ g ] [ 0 ]
18 %
%
%              n  p  1  1
20 %              t  >= 0
%              y  >= 0
22 %              ||f||_2 <= g

24 n = length(r);
p = size(G,1);

26
prob.a = sparse([ r'-rf, zeros(1, 2+p);
28              G, -eye(p), zeros(p,2);
                ones(1,n), zeros(1,1+p), -sum(w0) ]);
30 prob.c = [ zeros(n,1); zeros(p,1); 1; 0];
prob.cones{1}.type = 'MSK_CT_QUAD';
32 prob.cones{1}.sub = [n+p+1,n+1:n+p];
b = [ 1; zeros(p,1); 0 ];
34 prob.blc = b;
prob.buc = b;
36 prob.blx = [zeros(n,1); -inf*ones(p+1,1); 0];

38 [err, res]=mosekopt('minimize',prob);
y = res.sol.itr.xx(1:n);
40 f = res.sol.itr.xx(n+[1:p]);
g = res.sol.itr.xx(n+p+1);
42 t = res.sol.itr.xx(n+p+2);
x = y/t - w0;
44 risk = g/t;
ret = r'*(x+w0);

```

Listing 2: Largest Sharpe ratio without short-selling.

References

- [1] Richard C. Grinold and Ronald N. Kahn. *Active portfolio management*. McGraw-Hill, New York, 2 edition, 2000.
- [2] E. D. Andersen and Mosek employees. *The MOSEK optimization tools version 2.5. User's manual and reference*. MOSEK ApS, Fruebjergvej 3, Boks 16, 2100 Copenhagen O, 2002.
- [3] A. Ben-Tal and A Nemirovski. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. MPS/SIAM Series on Optimization. SIAM, 2001.
- [4] S.P. Boyd, S.J. Kim, L. Vandenberghe, and A. Hassibi. A Tutorial on Geometric Programming. Technical report, ISL, Electrical Engineering Department, Stanford University, Stanford, CA, 2004. Available at http://www.stanford.edu/~boyd/gp_tutorial.html.
- [5] Google Finance. Historical stock data. <http://www.google.com/finance/>.
- [6] Wikipedia. List of S&P 500 companies. http://en.wikipedia.org/wiki/List_of_S%26P_500_companies.