

# Solutions to Exercises

## Portfolio Optimization: Theory and Application Chapter 11 – Risk Parity Portfolios

Daniel P. Palomar (2025). *Portfolio Optimization: Theory and Application*.  
Cambridge University Press.

[portfoliooptimizationbook.com](http://portfoliooptimizationbook.com)

### Exercise 11.1: Change of variable

Show why  $\Sigma \mathbf{x} = \mathbf{b}/\mathbf{x}$  can be equivalently solved as  $\mathbf{C}\mathbf{x} = \mathbf{b}/\mathbf{x}$ , where  $\mathbf{C}$  is the correlation matrix defined as  $\mathbf{C} = \mathbf{D}^{-1/2}\Sigma\mathbf{D}^{-1/2}$  with  $\mathbf{D}$  a diagonal matrix containing  $\text{diag}(\Sigma)$  along the main diagonal. Would it be possible to use instead  $\mathbf{C} = \mathbf{M}^{-1/2}\Sigma\mathbf{M}^{-1/2}$ , where  $\mathbf{M}$  is not necessarily a diagonal matrix?

### Solution

Start by writing  $\Sigma \mathbf{x} = \mathbf{b}/\mathbf{x}$  as

$$\mathbf{D}^{-1/2}\Sigma\mathbf{D}^{-1/2}\tilde{\mathbf{x}} = \mathbf{D}^{-1/2}\mathbf{b}/(\mathbf{D}^{-1/2}\tilde{\mathbf{x}}) = \mathbf{b}/\tilde{\mathbf{x}},$$

where  $\tilde{\mathbf{x}} = \mathbf{D}^{1/2}\mathbf{x}$ . This leads to

$$\mathbf{C}\tilde{\mathbf{x}} = \mathbf{b}/\tilde{\mathbf{x}},$$

and we can recover  $\mathbf{x}$  as  $\mathbf{x} = \mathbf{D}^{-1/2}\tilde{\mathbf{x}} = \tilde{\mathbf{x}}/\sigma$ , where  $\sigma$  denote the volatilities, i.e., the diagonal elements of  $\mathbf{D}^{1/2}$ .

Now, to see if it would be possible to use instead  $\mathbf{C} = \mathbf{M}^{-1/2}\Sigma\mathbf{M}^{-1/2}$ , where  $\mathbf{M}$  is not necessarily a diagonal matrix, let's proceed similarly:

$$\mathbf{M}^{-1/2}\Sigma\mathbf{M}^{-1/2}\tilde{\mathbf{x}} = \mathbf{M}^{-1/2}\mathbf{b}/(\mathbf{M}^{-1/2}\tilde{\mathbf{x}}),$$

where  $\tilde{\mathbf{x}} = \mathbf{M}^{1/2}\mathbf{x}$ . The issue here is that

$$\mathbf{M}^{-1/2}\mathbf{b}/(\mathbf{M}^{-1/2}\tilde{\mathbf{x}}) \neq \mathbf{b}/\tilde{\mathbf{x}}$$

because the matrix  $\mathbf{M}$  is not diagonal, so it seems that a nondiagonal matrix cannot be used.

**Exercise 11.2:** Naive diagonal risk parity portfolio

If the covariance matrix is diagonal,  $\Sigma = D$ , then the system of nonlinear equations  $\Sigma \mathbf{x} = \mathbf{b}/\mathbf{x}$  has the closed-form solution  $\mathbf{x} = \sqrt{\mathbf{b}/\text{diag}(D)}$ . Explore whether a closed-form solution can be obtained for the rank-one plus diagonal case  $\Sigma = \mathbf{u}\mathbf{u}^\top + D$ .

**Solution**

If  $\Sigma = D = \text{Diag}(\mathbf{d})$ , then  $\Sigma \mathbf{x} = \mathbf{b}/\mathbf{x}$  can be written as

$$D\mathbf{x} = \text{Diag}(\mathbf{d})\mathbf{x} = \mathbf{d} \odot \mathbf{x} = \mathbf{b}/\mathbf{x}$$

which leads to

$$\mathbf{x}^2 = \mathbf{b}/\mathbf{d}$$

or

$$\mathbf{x} = \sqrt{\mathbf{b}/\mathbf{d}}.$$

Now, to see whether a closed-form solution can still be obtained for the rank-one plus diagonal case  $\Sigma = \mathbf{u}\mathbf{u}^\top + D$ , let's proceed similarly:

$$(\mathbf{u}\mathbf{u}^\top + D)\mathbf{x} = (\mathbf{u}^\top \mathbf{x})\mathbf{u} + \mathbf{d} \odot \mathbf{x} = \mathbf{b}/\mathbf{x}.$$

This leads to

$$(\mathbf{u}^\top \mathbf{x})(\mathbf{u}/\mathbf{d}) \odot \mathbf{x} + \mathbf{x}^2 = \mathbf{b}/\mathbf{d},$$

which does not seem to simplify as before. However, this can still be solved with the closed-form solution to a second-order equation:

$$\mathbf{x}^2 + (\mathbf{u}^\top \mathbf{x})(\mathbf{u}/\mathbf{d}) \odot \mathbf{x} - \mathbf{b}/\mathbf{d} = \mathbf{0}$$

with positive solution given by

$$x_i = \frac{-(\mathbf{u}^\top \mathbf{x})(u_i/d_i) + \sqrt{((\mathbf{u}^\top \mathbf{x})(u_i/d_i))^2 + 4b_i/d_i}}{2}, \quad i = 1, \dots, n$$

or, in vector form,

$$\mathbf{x} = \frac{-(\mathbf{u}^\top \mathbf{x})(\mathbf{u}/\mathbf{d}) + \sqrt{((\mathbf{u}^\top \mathbf{x})(\mathbf{u}/\mathbf{d}))^2 + 4(\mathbf{b}/\mathbf{d})}}{2}.$$

**Exercise 11.3:** Vanilla convex risk parity portfolio

The solution to the formulation

$$\begin{aligned} & \underset{\mathbf{x} \geq \mathbf{0}}{\text{maximize}} && \mathbf{b}^\top \log(\mathbf{x}) \\ & \text{subject to} && \sqrt{\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x}} \leq \sigma_0 \end{aligned}$$

is

$$\lambda \boldsymbol{\Sigma} \mathbf{x} = \mathbf{b}/\mathbf{x} \times \sqrt{\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x}}.$$

Can you solve for  $\lambda$  and rewrite the solution in a more compact way without  $\lambda$ ?

**Solution**

We can left-multiply both sides of  $\lambda \boldsymbol{\Sigma} \mathbf{x} / \sqrt{\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x}} = \mathbf{b}/\mathbf{x}$  by  $\mathbf{x}^\top$  to obtain:

$$\lambda \sqrt{\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x}} = \mathbf{x}^\top (\mathbf{b}/\mathbf{x}) = \mathbf{1}^\top \mathbf{b}.$$

Noting that at an optimal point it must be that the constraint is satisfied with equality:  $\sqrt{\mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x}} = \sigma_0$  (otherwise the objective value could be further increased), we can simplify it to

$$\lambda \sigma_0 = \mathbf{1}^\top \mathbf{b},$$

which leads to  $\lambda = \mathbf{1}^\top \mathbf{b} / \sigma_0$ . Then we can finally write the solution as

$$\frac{\mathbf{1}^\top \mathbf{b}}{\sigma_0^2} \times \boldsymbol{\Sigma} \mathbf{x} = \mathbf{b}/\mathbf{x}.$$

**Exercise 11.4:** Newton's method

Newton's method requires computing the direction  $\mathbf{d} = \mathbf{H}^{-1} \nabla f$  or, equivalently, solving the system of linear equations  $\mathbf{H} \mathbf{d} = \nabla f$  for  $\mathbf{d}$ . Explore whether a more efficient solution is possible by exploiting the structure of the gradient and Hessian:

$$\begin{aligned} \nabla f &= \boldsymbol{\Sigma} \mathbf{x} - \mathbf{b}/\mathbf{x}, \\ \mathbf{H} &= \boldsymbol{\Sigma} + \text{Diag}(\mathbf{b}/\mathbf{x}^2). \end{aligned}$$

**Solution**

Solving the system of linear equations  $\mathbf{H} \mathbf{d} = \nabla f$  for  $\mathbf{d}$  has a computational cost of  $O(n^3)$ . We will now assume  $\boldsymbol{\Sigma}^{-1}$  has been precomputed and then the resolution will have a cost of  $O(n^2)$ . The key is to apply the matrix inversion lemma to the Hessian matrix  $\mathbf{H} = \boldsymbol{\Sigma} + \mathbf{D}$ , where  $\mathbf{D} = \text{Diag}(\mathbf{b}/\mathbf{x}^2)$ ,

as follows:

$$H^{-1} = (\Sigma + D)^{-1} = \Sigma^{-1} - \Sigma^{-1}(I + D\Sigma^{-1})^{-1}D\Sigma^{-1}.$$

In particular, the steps are the following:

1. Compute preliminary term

$$v = \Sigma^{-1}\nabla f = \Sigma^{-1}(\Sigma x - b/x) = x - \Sigma^{-1}(b/x).$$

2. Form the matrix

$$M = I + D\Sigma^{-1}.$$

3. Solve for  $z$ :

$$Mz = Dv.$$

4. Compute final direction

$$d = v - \Sigma^{-1}z.$$

#### Exercise 11.5: MM algorithm

The MM algorithm requires the computation of the largest eigenvalue  $\lambda_{\max}$  of matrix  $\Sigma$ , which can be obtained from the eigenvalue decomposition of the matrix. A more efficient alternative is the *power iteration method*. Program both methods and compare their computational complexity.

#### Solution

First, we generate the covariance matrix:

```
library(microbenchmark)

# Generate covariance matrix
set.seed(42)
n <- 100
A <- matrix(rnorm(n^2), n, n)
Sigma <- t(A) %*% A
```

Then, we compute the maximum eigenvalue with the built-in function `eigen()`:

```

direct_nanoseconds = microbenchmark({
  lmd_max <- max(eigen(Sigma)$values)
}, unit = "nanoseconds", times = 100L)$time |> median()

cat(direct_nanoseconds, "nanoseconds used by the built-in function eigen()
  to compute the maximum eigenvalue of", lmd_max)

```

```

2743492 nanoseconds used by the built-in function eigen()
  to compute the maximum eigenvalue of 377.2282

```

Finally, we employ 20 iterations of the power iteration method (the number of iterations depends on the accuracy desired):

```

x0 <- rnorm(n)
power_iteration_nanoseconds = microbenchmark({
  u <- x0; for (i in 1:20) u <- Sigma %*% u
  lmd_max <- as.numeric(t(u) %*% Sigma %*% u / sum(u^2))
}, unit = "nanoseconds", times = 100L)$time |> median()

cat(power_iteration_nanoseconds, "nanoseconds used by the power iteration method
  to compute the maximum eigenvalue of", lmd_max)

```

```

1675598 nanoseconds used by the power iteration method
  to compute the maximum eigenvalue of 375.2165

```

### Exercise 11.6: Coordinate descent vs. SCA methods

Consider the vanilla convex formulation

$$\underset{\mathbf{x} \geq \mathbf{0}}{\text{minimize}} \quad \frac{1}{2} \mathbf{x}^\top \boldsymbol{\Sigma} \mathbf{x} - \mathbf{b}^\top \log(\mathbf{x}).$$

Implement the cyclical coordinate descent method and the parallel SCA method in a high-level programming language (e.g., R, Python, Julia, or MATLAB) and compare the convergence against the CPU time for these two methods. Then, re-implement these two methods in a low-level programming language (e.g., C, C++, C#, or Rust) and compare the convergence again. Comment on the difference observed.

## Solution

Let's construct a covariance matrix from stock market data:

```
library(xts)
library(pob)          # Market data used in the book
library(riskParityPortfolio)

# Prep data
N <- 200
Sigma <- cov(diff(log(SP500_2015to2020$stocks[, 1:N]))[-1])
sigma <- sqrt(diag(Sigma))
C <- cov2cor(Sigma)
b <- rep(1/N, N)
w_opt <- riskParityPortfolio(Sigma, b = b)$w
x_opt <- w_opt / as.vector(sqrt(w_opt %*% Sigma %*% w_opt))
opt_value <- 0.5 * x_opt %*% Sigma %*% x_opt - b %*% log(x_opt)

num_iter <- 10L
num_times <- 10L # to compute the cpu time
```

We can start with the cyclical coordinate descent method for the function  $f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^\top \mathbf{\Sigma} \mathbf{x} - \mathbf{b}^\top \log(\mathbf{x})$ . The elementwise minimization becomes

$$\underset{x_i \geq 0}{\text{minimize}} \quad \frac{1}{2}x_i^2 \Sigma_{ii} + x_i(\mathbf{x}_{-i}^\top \mathbf{\Sigma}_{-i,i}) - b_i \log x_i,$$

where  $\mathbf{x}_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N)$  denotes the variable  $\mathbf{x}$  without the  $i$ th element and  $\mathbf{\Sigma}_{-i,i}$  denotes the  $i$ th column of matrix  $\mathbf{\Sigma}$  without the  $i$ th element. Setting the partial derivative with respect to  $x_i$  to zero gives us the second-order equation

$$\Sigma_{ii}x_i^2 + (\mathbf{x}_{-i}^\top \mathbf{\Sigma}_{-i,i})x_i - b_i = 0$$

with positive solution given by

$$x_i = \frac{-\mathbf{x}_{-i}^\top \mathbf{\Sigma}_{-i,i} + \sqrt{(\mathbf{x}_{-i}^\top \mathbf{\Sigma}_{-i,i})^2 + 4\Sigma_{ii}b_i}}{2\Sigma_{ii}}.$$

```

library(microbenchmark)
library(dplyr)

#
# Cyclical Spinu coordinate descent algorithm
#
x <- sqrt(b)/sqrt(rowSums(Sigma))
df <- data.frame(
  "k" = 0L,
  "cpu time k" = 0,
  "obj_value" = 0.5 * x %**% Sigma %**% x - b %**% log(x),
  "gap" = 0.5 * x %**% Sigma %**% x - b %**% log(x) - opt_value,
  "method" = "Cyclical coordinate descent",
  check.names = FALSE
)

for (k in 1:num_iter) {
  cpu_time <- microbenchmark({
    x_new <- x
    for (i in 1:N) {
      Sigma_xk_i <- as.numeric(x_new[-i] %**% Sigma[-i, i])
      x_new[i] <- (- Sigma_xk_i + sqrt(Sigma_xk_i^2 + 4*Sigma[i, i]*b[i]))/(2*Sigma[i, i])
    }
  }, unit = "microseconds", times = num_times)$time |> median()
  x <- as.numeric(x_new)

  df <- rbind(df, list(
    "k" = k,
    "cpu time k" = cpu_time,
    "obj_value" = 0.5 * x %**% Sigma %**% x - b %**% log(x),
    "gap" = 0.5 * x %**% Sigma %**% x - b %**% log(x) - opt_value,
    "method" = "Cyclical coordinate descent"))
}

```

The SCA method obtains the iterates  $\mathbf{x}^0, \mathbf{x}^1, \mathbf{x}^2, \dots$  by solving a sequence of simpler surrogate problems. In particular, the following surrogate can be used for the term  $\mathbf{x}^\top \Sigma \mathbf{x}$  around the current point  $\mathbf{x} = \mathbf{x}^k$ :

$$\frac{1}{2} \mathbf{x}^\top \Sigma \mathbf{x} \approx \frac{1}{2} (\mathbf{x}^k)^\top \Sigma \mathbf{x}^k + (\Sigma \mathbf{x}^k)^\top (\mathbf{x} - \mathbf{x}^k) + \frac{1}{2} (\mathbf{x} - \mathbf{x}^k)^\top \text{Diag}(\Sigma) (\mathbf{x} - \mathbf{x}^k),$$

where  $\text{Diag}(\Sigma)$  is a diagonal matrix containing the diagonal of  $\Sigma$ . We can now solve our original problems by solving instead a sequence of surrogate problems

$$\underset{\mathbf{x} \geq \mathbf{0}}{\text{minimize}} \quad \frac{1}{2} \mathbf{x}^\top \text{Diag}(\Sigma) \mathbf{x} + \mathbf{x}^\top (\Sigma - \text{Diag}(\Sigma)) \mathbf{x}^k - \mathbf{b}^\top \log(\mathbf{x}),$$

from which setting the gradient to zero gives the second-order equation

$$\Sigma_{ii} x_i^2 + ((\Sigma - \text{Diag}(\Sigma)) \mathbf{x}^k)_i x_i - b_i = 0$$

with positive solution

$$x_i = \frac{-((\Sigma - \text{Diag}(\Sigma))\mathbf{x}^k)_i + \sqrt{((\Sigma - \text{Diag}(\Sigma))\mathbf{x}^k)_i^2 + 4\Sigma_{ii}b_i}}{2\Sigma_{ii}}$$

```
#
# Parallel Spinu SCA
#
x <- sqrt(b)/sqrt(rowSums(Sigma))
gamma <- 1
eps <- 0.1
df <- rbind(df, list(
  "k"          = 0L,
  "cpu time k" = 0,
  "obj_value"  = 0.5 * x %>% Sigma %>% x - b %>% log(x),
  "gap"        = 0.5 * x %>% Sigma %>% x - b %>% log(x) - opt_value,
  "method"     = "SCA")
)

Sigma_Diag_Sigma <- Sigma - diag(diag(Sigma))
for (k in 1:num_iter) {
  cpu_time <- microbenchmark({
    Sigma_Diag_Sigma_xk <- Sigma_Diag_Sigma %>% x
    x_hat <- (-Sigma_Diag_Sigma_xk + sqrt(Sigma_Diag_Sigma_xk^2 + 4*diag(Sigma)*b))/(2*diag(Sigma))
    x_new <- gamma*x_hat + (1 - gamma)*x
  }, unit = "microseconds", times = num_times)$time |> median()
  x <- as.numeric(x_new)
  gamma <- gamma * (1 - eps*gamma)

  df <- rbind(df, list(
    "k"          = k,
    "cpu time k" = cpu_time,
    "obj_value"  = 0.5 * x %>% Sigma %>% x - b %>% log(x),
    "gap"        = 0.5 * x %>% Sigma %>% x - b %>% log(x) - opt_value,
    "method"     = "SCA"))
}
```

Plot convergence:




```

library(ggplot2)
library(dplyr)
library(scales)

# Compute cumulative CPU time over iterations
df <- df |>
  group_by(method) |>
  mutate("CPU time [ms]" = cumsum(`cpu time k`)/1e6) |>
  ungroup()

# Plots
df |>
  ggplot(aes(x = k, y = gap, color = method)) +
  geom_line(linewidth = 1.2) +
  scale_y_log10(breaks = trans_breaks("log10", function(x) 10^x),
               labels = trans_format("log10", math_format(10^.x))) +
  ggtitle("Optimality gap vs. iterations")

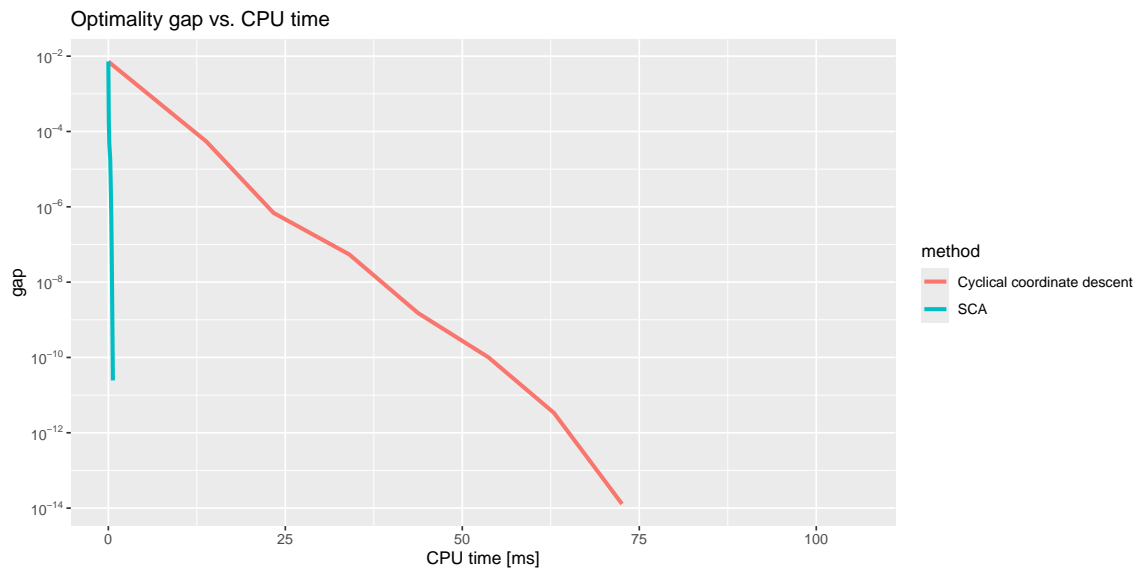
```



```

df |>
  ggplot(aes(x = `CPU time [ms]`, y = gap, color = method)) +
  geom_line(linewidth = 1.2) +
  scale_y_log10(breaks = trans_breaks("log10", function(x) 10^x),
               labels = trans_format("log10", math_format(10^.x))) +
  ggtitle("Optimality gap vs. CPU time")

```



We can observe the much faster convergence of the parallel SCA method compared to the sequential cyclical coordinate descent.

We leave the implementation in a low-level programming language (e.g., C, C++, C#, or Rust) to the user to observe whether the difference in convergence speed reduces or remains the same.