The Markowitz Portfolio

SAPIENZA UNIVERSITÀ DI ROMA

Outline

Managerial motivation

Guided construction of the model

Worked example by hand

Case study

Results

Sensitivity analysis

Exercises

Learning objectives (1/2)

At the end of this chapter you will be able to:

- formulate and solve a convex QP;
- build and read the efficient frontier;

Learning objectives (2/2)

At the end of this chapter you will be able to:

- understand the effect of mandate constraints;
- be wary of the estimates of the expected returns.

Prerequisites

- From this module, Chapter 2 «Background: convexity, quadratic programming and KKT conditions»: convex quadratic programs, the KKT conditions and the multipliers as shadow prices
- From the module *Convex Optimization*, Chapter 4 «Quadratic functions»: quadratic forms and covariance matrices as positive semidefinite matrices

Managerial motivation

Introduction: Managerial motivation

- How much should we invest in each asset in order to balance **expected return** and **risk**?
- The Markowitz model is the **most famous QP in history**: it minimizes the variance of the portfolio subject to a minimum return
- Class of the problem: **convex QP**

The efficient frontier

- The result is not a single number but an **efficient frontier**
- It is the complete menu of risk-return trade-offs
- The decision maker chooses on this frontier according to their own degree of **risk aversion**

A structure that reappears everywhere

- The problem was born in finance (Markowitz, 1952, Nobel prize)
- Its structure — **allocating a budget among risky and correlated alternatives** — reappears everywhere:
  - mixes of R&D projects
  - customer portfolios
  - allocation of production capacity across volatile markets
- It is the natural bridge between optimization and statistics: the data of the model are **estimates**, with all the consequences that this entails

Summary

- Markowitz: minimize the risk (variance) while guaranteeing a minimum return
- The result is an **efficient frontier** of risk-return trade-offs
- The data are **estimates**: optimization meets statistics

Guided construction of the model

Introduction: Guided construction of the model

- We describe the problem in words, then we fix **data** and **variables**
- We write the **minimum-variance QP model** and explain objective and constraints
- We see two **equivalent formulations** and why the QP is **convex**

The problem in words

- **We decide** the share of capital to invest in each asset
- **The objective** is to minimize the risk of the portfolio, measured by the variance of its return
- **The constraints**:
  - invest the whole capital (the shares sum to 1)
  - guarantee a minimum expected return
  - respect the minimum and maximum share of each asset

Data (input of the model) (1/2)

Symbol	Type	Meaning
n	$\in \mathbb{Z}_{\geq 1}$	number of assets; the assets are indexed by $i \in \{1, 2, \ldots, n\}$
μ_i	$\in \mathbb{Q}$	expected return of asset i (per year)
q_{ij}	$\in \mathbb{Q}$	covariance between the returns of assets i and j; the matrix $\boldsymbol{Q} = (q_{ij})$ is positive semidefinite ($\boldsymbol{Q} \succeq 0$)

Data (input of the model) (2/2)

Symbol	Type	Meaning
$\bar{r}$	$\in \mathbb{Q}$	minimum required expected return
ℓ_i, u_i	$\in \mathbb{Q}$, $0 \leq \ell_i \leq u_i \leq 1$	minimum and maximum share that can be invested in asset i

Decision variables

- We introduce the following n non-negative variables:

  x_i = share of capital invested in asset i, $\quad \forall i \in \{1, 2, \ldots, n\}$

- Using these variables, we write a QP model for the problem

QP model

Minimum-variance QP model with a minimum return

$$\min \quad \sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij}\, x_i\, x_j \tag{6.1a}$$

$$\text{subject to} \quad \sum_{i=1}^{n} \mu_i\, x_i \geq \bar{r}, \tag{6.1b}$$

$$\sum_{i=1}^{n} x_i = 1, \tag{6.1c}$$

$$x_i \leq u_i, \quad \forall i \in \{1, 2, \ldots, n\}, \tag{6.1d}$$

$$x_i \geq \ell_i, \quad \forall i \in \{1, 2, \ldots, n\}. \tag{6.1e}$$

The objective function

- The **quadratic** objective function minimizes the variance of the return of the portfolio
- It contains the covariances q_{ij}, and this is where **diversification** comes from
- Combining weakly correlated assets lowers the overall risk **below that of the individual components**

Minimum return and budget

- The linear constraint $\sum_i \mu_i x_i \geq \bar{r}$ imposes that the expected return of the portfolio reaches the threshold $\bar{r}$ (one linear constraint)
  - it is the "**dial**" with which the efficient frontier is traced out
  - its multiplier tells us how much variance each additional point of return costs
- The linear constraint $\sum_i x_i = 1$ imposes that the whole capital is invested (one linear constraint)

Bounds per asset

- The linear constraints $x_i \leq u_i$ are the **caps per asset**, typical regulatory or mandate constraints (n linear constraints)
- The constraints $x_i \geq \ell_i$ define the variables of the model
- With $\ell_i = 0$: **short selling is forbidden**

Equivalent formulations

Mean-variance: the risk goes into the objective with a weight λ

$$\max \sum_{i=1}^{n} \mu_i x_i - \lambda \sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij} x_i x_j$$

Maximum return with bounded risk: the risk goes into the constraints

$$\max \sum_{i=1}^{n} \mu_i x_i \quad \text{with variance} \leq \bar{\sigma}^2$$

By varying $\bar{r}$ (or λ, or $\bar{\sigma}$) **the same frontier** is traced out.

Convexity

- Every covariance matrix is **positive semidefinite**
- Hence the QP is **convex**
- The optimum that is found is a **certified global** one

Summary

- Variables x_i: shares of capital invested in the assets
- Quadratic objective (variance); linear constraints: minimum return, budget, bounds per asset
- Three equivalent formulations (minimum variance, mean-variance, bounded risk) generate the same frontier
- $\boldsymbol{Q} \succeq 0$: the QP is convex, the optimum is global

Worked example by hand

Introduction: Worked example by hand

- Only two assets, uncorrelated
- We compute by hand the **minimum-variance** portfolio
- We will see the **diversification effect** in its purest form

Example: Two uncorrelated assets (1/2)

Volatilities $\sigma_1 = 20\%$, $\sigma_2 = 30\%$, zero correlation.
Minimize the variance with $x_1 + x_2 = 1$ (without the return constraint).

Step 1. Substituting $x_2 = 1 - x_1$:

$$f(x_1) = 0.04\, x_1^2 + 0.09\,(1 - x_1)^2$$

Step 2.

$$f'(x_1) = 0.08\, x_1 - 0.18\,(1 - x_1) = 0$$

$$\Rightarrow \; x_1 = \frac{0.09}{0.04 + 0.09} = \frac{9}{13} \approx 69.2\%$$

Example: Two uncorrelated assets (2/2)

Step 3. Variance of the portfolio:

$$f(9/13) = 0.04 \cdot 0.479 + 0.09 \cdot 0.095 = 0.0277$$

that is a volatility of 16.6%

- The portfolio is **less risky than either of the two assets**
- This is the diversification effect in its purest form: the optimal weights are **inversely proportional to the variances**

Summary

- With two uncorrelated assets the minimum is found by setting the derivative to zero
- $x_1 = 9/13$: volatility 16.6%, lower than that of either asset
- Optimal weights inversely proportional to the variances

Case study

Introduction: Case study

- Eight sector ETFs with simulated monthly returns
- μ and $\boldsymbol{Q}$ are **estimated** from the data, as in practice
- We see the model written in **Python with Gurobi**

The data of the case study

- Eight **sector ETFs** (*Exchange Traded Fund*: listed funds that replicate a basket of securities of a given sector)
- 60 monthly returns simulated with a single-market-factor model
- File `markowitz_rendimenti.csv`
- $\boldsymbol{\mu}$ and $\boldsymbol{Q}$ are **estimated** from these data, as in practice

Annualized statistics

Asset	μ	vol	Asset	μ	vol
ENE	19.12%	18.65%	SAN	2.80%	11.93%
FIN	8.27%	19.59%	CON	11.25%	11.67%
TEC	1.58%	29.27%	UTL	11.99%	7.53%
IND	9.97%	16.60%	MAT	20.93%	23.93%

The model in Gurobi

```
m = gp.Model("markowitz")
# shares, 0 <= x_i <= u
x = m.addVars(n, ub=u, name="x")
m.addConstr(x.sum() == 1, name="budget")
m.addConstr(gp.quicksum(mu[i]*x[i]
                        for i in range(n)) >= r_min)
m.setObjective(gp.quicksum(Q[i, j]*x[i]*x[j]
                           for i in range(n)
                           for j in range(n)),
               GRB.MINIMIZE)
m.optimize()
```

Summary

- Eight sector ETFs, 60 months of simulated returns
- Expected returns and covariances are **estimates** from the data
- In Gurobi the quadratic objective is written directly with `quicksum` over the products $x_i x_j$

Results

Introduction: Results

- We compare the **minimum-variance** portfolio with the **equally weighted** $1/n$ one
- We draw the **efficient frontier**, with and without caps per asset
- We look at how the optimal **composition** changes along the frontier

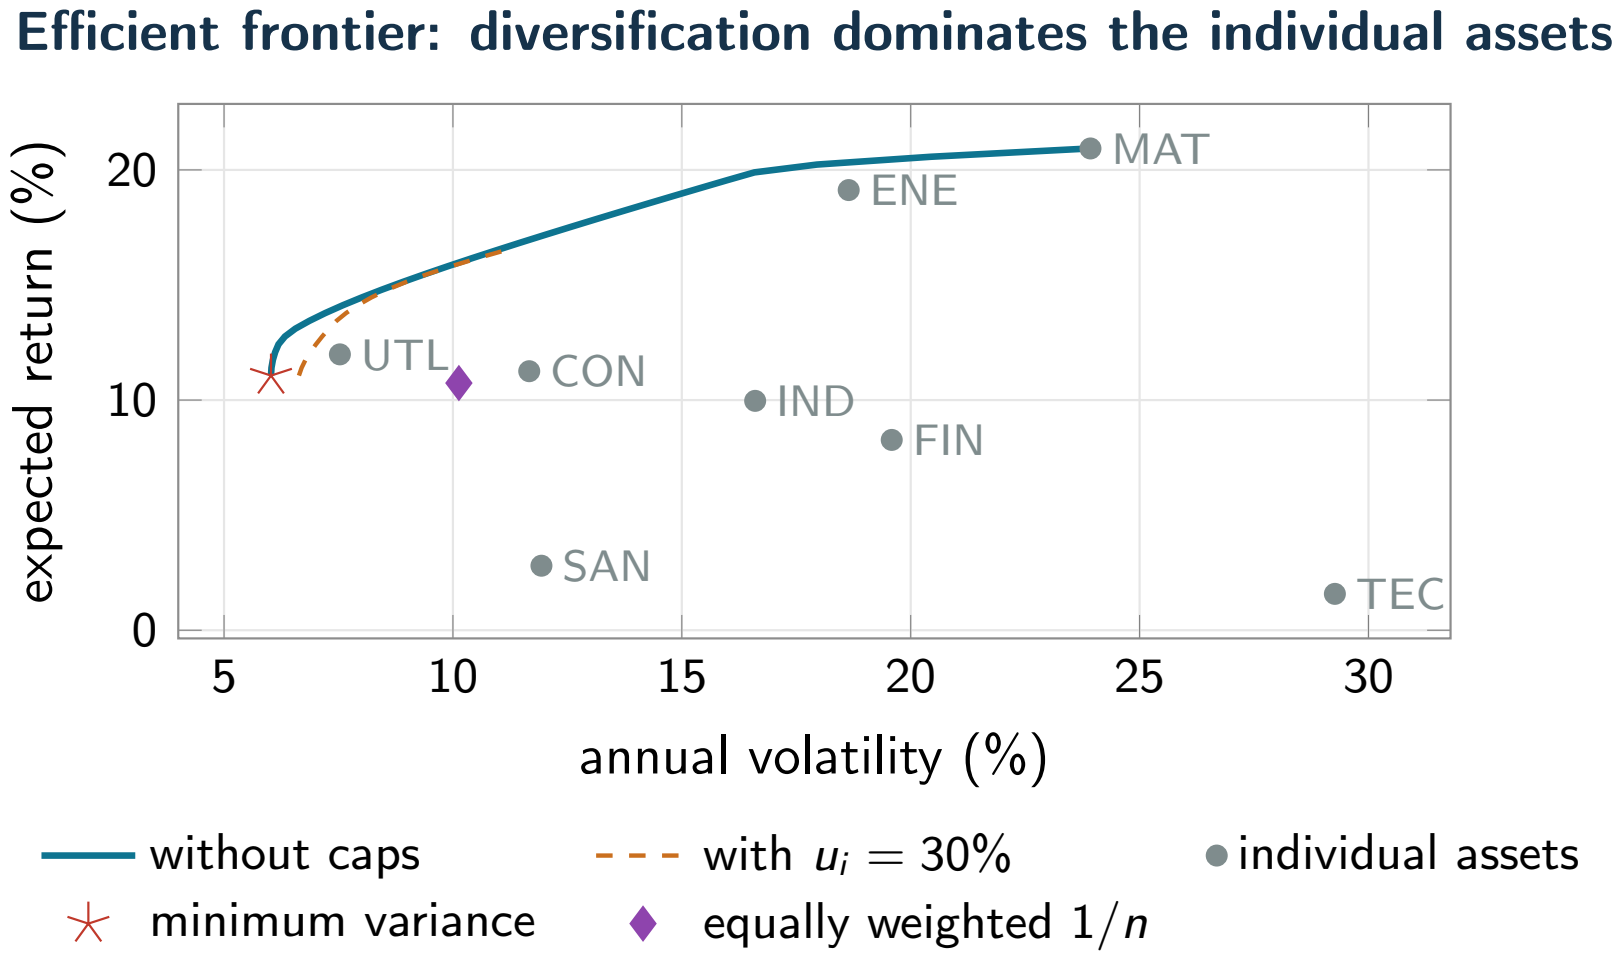
Minimum variance versus equally weighted

Portfolio	Return	Volatility
Global minimum variance	11.06%	6.03%
Equally weighted $(1/n)$	10.74%	10.13%

Minimum-variance composition:

ENE	IND	SAN	CON	UTL
5.2%	1.9%	11.5%	22.3%	58.3%

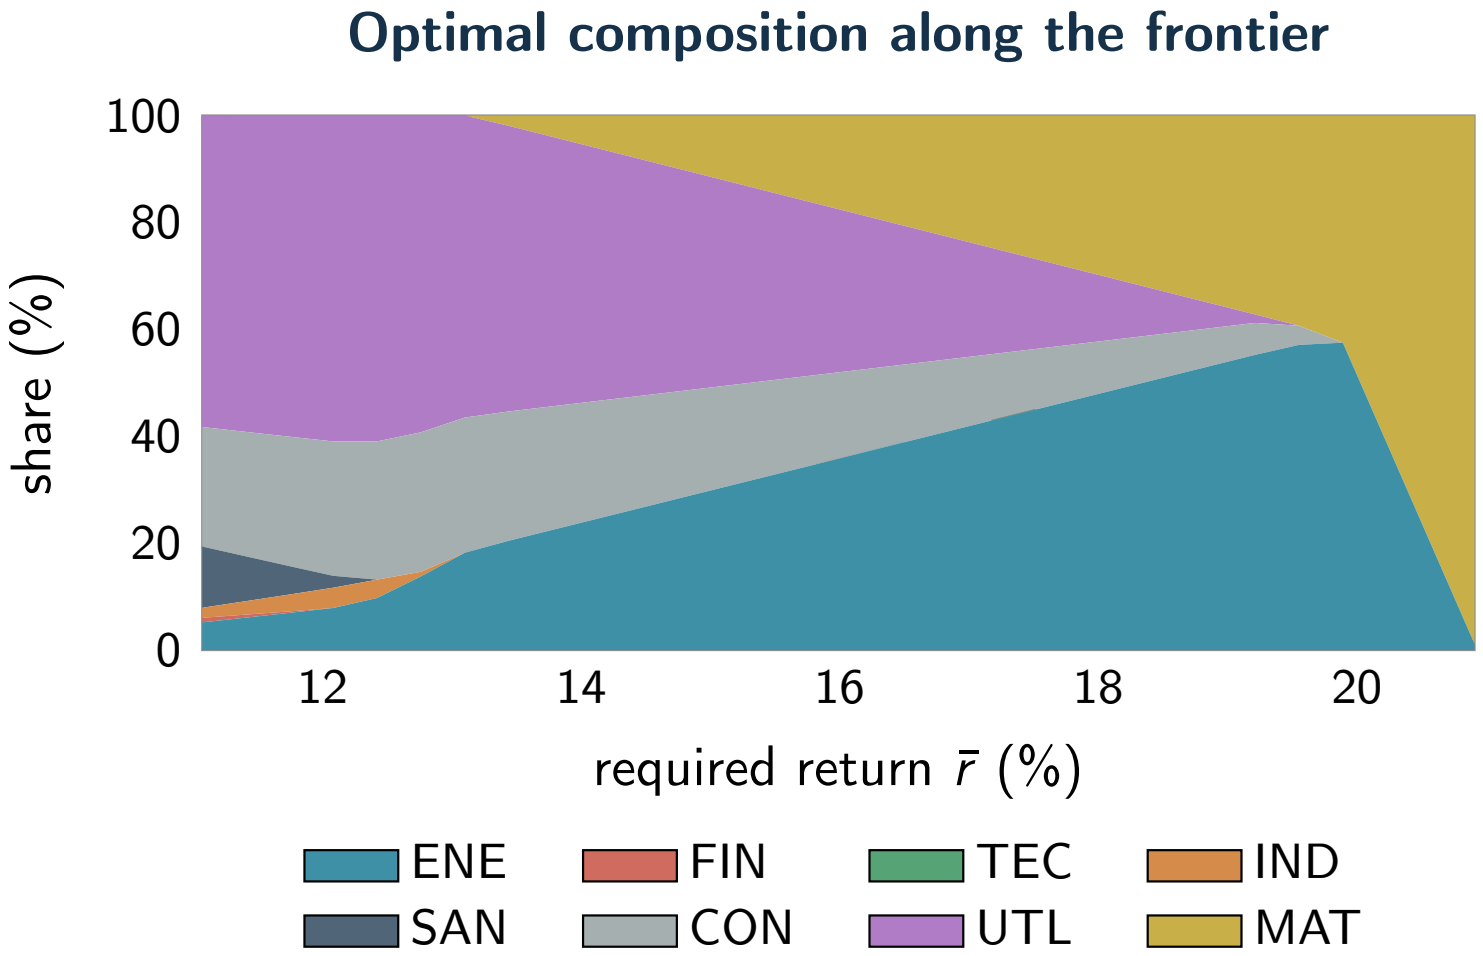
Efficient frontier



How to read the risk-return plane

- The eight assets: **grey dots**
- Efficient frontier without caps: **teal line**
- Frontier with a 30% cap per asset: **dashed orange**
- Minimum-variance portfolio: **red star**
- Equally weighted $1/n$: **purple diamond**, which turns out to be **dominated**

Optimal composition along the frontier



How to read the composition

- Stacked areas (in %) as the required return $\bar{r}$ grows
- On the left the **defensive mix** dominates (UTL, CON, SAN)
- Towards the right the portfolio concentrates progressively on the **high-return** assets (MAT, ENE)
- At the extreme: the **single-asset** portfolio

Three messages from the figure

1. The minimum-variance portfolio (volatility 6%) is much less risky than the best individual asset (UTL, 7.5%) while still returning 11%
2. The equally weighted $1/n$ portfolio — the "I know nothing" strategy — is clearly **dominated**: the same region of return but almost twice the volatility
3. The cap $u_i = 30\%$ cuts off the upper part of the frontier: the **cost of the mandate constraints** can be *seen* as the distance between the two curves

Summary

- Minimum variance: return 11.06%, volatility 6.03%
- The equally weighted portfolio is dominated: similar return, almost twice the volatility
- Along the frontier we move from the defensive mix to the high-return assets
- The mandate constraints have a visible cost on the frontier

Sensitivity analysis

Introduction: Sensitivity analysis

- We let the minimum required return $\bar{r}$ vary
- We will understand when the return constraint is **active** and how much it costs
- Finally we will see the **fragility of the estimates** of the expected returns

Volatility as $\bar{r}$ varies

$\bar{r}$	Volatility	
6%	6.03%	constraint **not** active: same as min variance
8%	6.03%	same
10%	6.03%	same
12%	6.10%	$d(\text{variance})/d(\bar{r}) \approx 0.0221$

A constraint that costs nothing

- Up to $\bar{r} = 11\%$ the return constraint is **inactive**
- The minimum-variance portfolio already returns 11.06%: asking for "at least 8%" costs nothing and the multiplier is zero
- Only beyond 11.06% does the constraint bite: every additional point of return is paid for in variance (multiplier ≈ 0.022 at $\bar{r} = 12\%$)
- Recognizing inactive constraints avoids "paying" for constraints that are free

The fragility of the estimates

Observation 1: The fragility of the estimates

In the simulated data the asset TEC has a true α of 11% per year, but over 60 months the estimated return is 1.6%: the noise dominates.

This is the most important practical lesson of Markowitz: ***the estimates of the expected returns are far more unstable than those of the covariances****, and optimized portfolios chase the estimation errors.*

Remedies used in practice

- The constraints u_i on the shares per asset
- **Shrinkage** of the estimates
- Optimizing risk only (**minimum variance**)

Summary

- Below 11.06% the return constraint is inactive: zero multiplier
- Beyond it, every point of return costs variance (≈ 0.022 at $\bar{r} = 12\%$)
- The estimates of the expected returns are fragile: optimization chases the estimation errors
- Remedies: caps, shrinkage, minimum variance

Exercises

Introduction: Exercises

- Five exercises, by hand and with the script of the case study
- They concern the correlation, the caps per asset, the stability of the estimates, the tracking error and the transaction costs

Exercise: check

- Verify by hand the example with two assets and $\rho = 0.5$: the formula becomes

$$x_1 = \frac{\sigma_2^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}$$

- Is diversification still worthwhile?

Exercise: frontier

- Recompute the frontier with $u_i = 20\%$
- How much maximum return is lost?
- Draw the three frontiers ($u = 1, 0.3, 0.2$) in the same plot

Exercise: stability of the estimates

- Run the script again with a different random seed (`default_rng(1)`)
- Compare the optimal compositions for the same $\bar{r}$: how much do they change?
- Which operational implication do you draw from this?

Exercise: tracking error

- The benchmark is the equally weighted portfolio $x_i^B = 1/8$
- Minimize the tracking error

$$\sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij}(x_i - x_i^B)(x_j - x_j^B)$$

  with expected return $\geq 12\%$
- How far must one depart from the benchmark?

Exercise: transaction costs

- Starting from the equally weighted portfolio, add the quadratic cost

$$\kappa \sum_{i=1}^{n} (x_i - 1/8)^2, \qquad \kappa = 0.5$$

- How does the composition change with respect to the "clean" optimum?

Summary

- Effect of the correlation on diversification
- Cost of the caps per asset on the frontier
- Stability of the compositions with respect to the data
- Variants of the model: tracking error with respect to a benchmark and quadratic transaction costs

Final summary (1/2)

Model:

- Convex QP: minimum variance with a minimum return, budget and bounds per asset
- Equivalent formulations: mean-variance, bounded risk

Diversification:

- Combining weakly correlated assets lowers the risk below that of the individual assets

Final summary (2/2)

Efficient frontier:

- Menu of the risk-return trade-offs; the equally weighted portfolio is dominated
- The mandate constraints have a visible cost

Estimates:

- Inactive constraints cost nothing
- The estimates of the expected returns are fragile: use caps, shrinkage or minimum variance