Operations Research Lab
Continuous optimization models for Management Engineering

Fabio Furini

Sapienza University of Rome

Academic year 2026–2027

Course outline

Tools

- Introduction and philosophy of the course
- Theory: linear programming
- Theory: nonlinear optimization
- Implementation: linear and nonlinear models

Deterministic models

- Production and inventory (LP/QP)
- Supply chain and CO_2 (LP/NLP)
- Markowitz portfolio (QP)
- Pricing (NLP)
- Advertising budget (convex NLP)
- Continuous location (convex NLP)
- EV charging (LP/QP)
- Queues and capacity (convex NLP)

Uncertainty and machine learning

- Newsvendor (stochastic LP)
- VaR and CVaR (LP)
- Arbitrage and pricing (LP)
- Support Vector Machine (QP)

Lab organization

Introduction

What this lab is about

- Turning business decisions into optimization models: how much to produce, where to locate, what price to set, how much risk to accept.
- Implementing the models in **Python + Gurobi**, solving them and — above all — questioning them: how much is one extra hour of capacity worth? does the solution hold if the data change by 5%?
- **Continuous variables** only: duality, shadow prices and KKT conditions apply — the tools that turn numbers into economic explanations.

The final question of every lab session

Not just *"what is the optimum?"*, but *"which decision do we recommend, and how robust is it?"*

Notation and abbreviations

- **LP** linear programming; **QP** quadratic; **NLP** nonlinear. *Convex* = every local minimum is global.
- Scalars and indices in lowercase (x_{it}, λ); **vectors** in lowercase bold ($\boldsymbol{x}$); **matrices** in uppercase bold ($\boldsymbol{Q}$); **numbered objects** with indices over enumerated sets: $i \in \{1, 2, \ldots, n\}$ — the notation of the text is already that of the model; rational data ($\mathbb{Q}$), integer counts ($\mathbb{Z}_{\geq 1}$).
- Dual variables π_i, reduced costs $\bar{c}_j$, slacks $\bar{s}_i$; the bar denotes the values of a feasible solution, the tilde those of an optimal one ($\tilde{x}_j$, $\tilde{z}$).
- In every model: first the data (with their type), then the variables, then objective and constraints explained line by line.

Theory: linear programming

The primal-dual pair (general form)

Data: costs c_j $(j \in N)$, coefficients a_{ij}, right-hand sides b_i $(i \in M)$. $M_{\leq}, M_{=}, M_{\geq}$ partition M by direction; $N_{\geq 0}, N_{\gtreqless 0}, N_{\leq 0}$ partition N by sign.

$$
\begin{aligned}
\text{(P)} \quad \max \ & \sum_{j \in N} c_j x_j \\
\text{subject to} \ & \sum_{j \in N} a_{ij} x_j \leq b_i, && \forall i \in M_{\leq}, \\
& \sum_{j \in N} a_{ij} x_j = b_i, && \forall i \in M_{=}, \\
& \sum_{j \in N} a_{ij} x_j \geq b_i, && \forall i \in M_{\geq}, \\
& x_j \geq 0, && \forall j \in N_{\geq 0}, \\
& x_j \gtreqless 0, && \forall j \in N_{\gtreqless 0}, \\
& x_j \leq 0, && \forall j \in N_{\leq 0}.
\end{aligned}
$$

$$
\begin{aligned}
\text{(D)} \quad \min \ & \sum_{i \in M} b_i \pi_i \\
\text{subject to} \ & \sum_{i \in M} a_{ij} \pi_i \geq c_j, && \forall j \in N_{\geq 0}, \\
& \sum_{i \in M} a_{ij} \pi_i = c_j, && \forall j \in N_{\gtreqless 0}, \\
& \sum_{i \in M} a_{ij} \pi_i \leq c_j, && \forall j \in N_{\leq 0}, \\
& \pi_i \geq 0, && \forall i \in M_{\leq}, \\
& \pi_i \gtreqless 0, && \forall i \in M_{=}, \\
& \pi_i \leq 0, && \forall i \in M_{\geq}.
\end{aligned}
$$

The rules of duality

Primal (max)		Dual (min)
constraint $\leq b_i$	$(i \in M_{\leq})$	variable $\pi_i \geq 0$
constraint $= b_i$	$(i \in M_{=})$	variable $\pi_i \gtreqless 0$
constraint $\geq b_i$	$(i \in M_{\geq})$	variable $\pi_i \leq 0$
variable $x_j \geq 0$	$(j \in N_{\geq 0})$	constraint $\geq c_j$
variable $x_j \gtreqless 0$	$(j \in N_{\gtreqless 0})$	constraint $= c_j$
variable $x_j \leq 0$	$(j \in N_{\leq 0})$	constraint $\leq c_j$

- Equalities in the primal give free dual variables (and vice versa).
- For a *minimization* primal the table is read from right to left.

The two duality theorems

Weak duality

For every $(\bar{x}_1, \ldots, \bar{x}_n)$ feasible for (P) and every $(\bar{\pi}_1, \ldots, \bar{\pi}_m)$ feasible for (D):

$$\sum_{j \in N} c_j \, \bar{x}_j \; \leq \; \sum_{i \in M} b_i \, \bar{\pi}_i$$

(any feasible dual is an upper bound for the primal).

Strong duality

(P) has an optimum $(\tilde{x}_1, \ldots, \tilde{x}_n) \iff$ (D) has an optimum $(\tilde{\pi}_1, \ldots, \tilde{\pi}_m)$, and then the gap closes:

$$\sum_{j \in N} c_j \, \tilde{x}_j \; = \; \sum_{i \in M} b_i \, \tilde{\pi}_i.$$

Building the dual with the rules: the 2×2 LP

Canonical maximization form: every $\leq$ constraint gives a $\pi_i \geq 0$; every variable ≥ 0 gives a constraint $\geq c_j$ (coefficients read by column); the objective swaps roles.

$$
\begin{array}{rrcrcr}
\max & 30\,x_1 & + & 50\,x_2 & & \\
\text{subject to} & x_1 & + & 3\,x_2 & \leq & 90, \\
& 2\,x_1 & + & x_2 & \leq & 80, \\
& x_1, & & x_2 & \geq & 0.
\end{array}
\qquad
\begin{array}{rrcrcr}
\min & 90\,\pi_1 & + & 80\,\pi_2 & & \\
\text{subject to} & \pi_1 & + & 2\,\pi_2 & \geq & 30, \\
& 3\,\pi_1 & + & \pi_2 & \geq & 50, \\
& \pi_1, & & \pi_2 & \geq & 0.
\end{array}
$$

Building the dual: all cases in a single LP

A *minimization* problem with three directions and three signs (table read from right to left):

$$
\begin{array}{rrcl}
\min & 5\,x_1 + 8\,x_2 - 9\,x_3 & & \\
\text{subj. to} & x_1 + x_2 & \geq & 30, \\
& x_1 + x_2 - x_3 & = & 100, \\
& x_1 - 2\,x_2 & \leq & -20, \\
& x_1 & \geq & 0, \\
& x_2 & \geq & 0, \\
& x_3 & \leq & 0.
\end{array}
$$

$$
\begin{array}{rrcl}
\max & 30\,\pi_1 + 100\,\pi_2 - 20\,\pi_3 & & \\
\text{subj. to} & \pi_1 + \pi_2 + \pi_3 & \leq & 5, \\
& \pi_1 + \pi_2 - 2\,\pi_3 & = & 8, \\
& - \pi_2 & \geq & -9, \\
& \pi_1 & \geq & 0, \\
& \pi_2 & \geq & 0, \\
& \pi_3 & \leq & 0.
\end{array}
$$

The LP optimality conditions: complementary slackness

Unknown variables plain; bar = values of a feasible solution; tilde = of an optimal one. Slack and reduced cost (the slack of the *dual* constraint; the bar in $\bar{c}_j$ is part of the name):

$$s_i = \sum_{j \in N} a_{ij} x_j - b_i, \qquad \bar{c}_j = c_j - \sum_{i \in M} a_{ij} \pi_i.$$

Theorem (complementary slackness)

Feasible $(\bar{x}_1, \ldots, \bar{x}_n)$, $(\bar{\pi}_1, \ldots, \bar{\pi}_m)$ are optimal for (P) and (D) $\iff$

$$\bar{\pi}_i \cdot \bar{s}_i = 0 \;\; \forall i \in M, \qquad \bar{x}_j \cdot \bar{c}_j = 0 \;\; \forall j \in N.$$

- Slack constraint $\Rightarrow$ zero dual; variable $\neq 0 \Rightarrow$ zero reduced cost (implications: degeneracy).

The signs of the shadow prices

One convention: $\tilde{\pi}_i = \dfrac{\partial \tilde{z}}{\partial b_i}$, valid for $b_i \in [b_i^{\min}, b_i^{\max}]$ (validity range)

- Two questions: does increasing b_i enlarge ($\leq$) or shrink ($\geq$) the feasible region? And a larger region can never worsen the optimum.

Constraint direction	minimization	maximization
$\leq$ ($b_i \uparrow \Rightarrow$ larger region)	$\tilde{\pi}_i \leq 0$	$\tilde{\pi}_i \geq 0$
$\geq$ ($b_i \uparrow \Rightarrow$ smaller region)	$\tilde{\pi}_i \geq 0$	$\tilde{\pi}_i \leq 0$
$=$	any sign (free dual)	

- In a minimization: $\leq$ (capacity) $\Rightarrow$ loosening reduces the cost, $|\tilde{\pi}_i| =$ marginal saving; $\geq$ (requirement) $\Rightarrow$ $\tilde{\pi}_i =$ marginal cost. Non-binding constraint: $\tilde{\pi}_i = 0$.

Duality on the 2×2 LP: solving gives, we verify

1. Solving gives: $(\tilde{x}_1, \tilde{x}_2) = (30, 20)$, $\tilde{z} = 1900$, $(\tilde{\pi}_1, \tilde{\pi}_2) = (14, 8)$, $(\bar{c}_1, \bar{c}_2) = (0, 0)$; both constraints binding.
2. Check (complementary slackness): $x_1, x_2 \neq 0 \Rightarrow$ dual constraints binding: $14 + 16 = 30 = c_1$, $42 + 8 = 50 = c_2$. ✓
3. Check of strong duality: $90 \cdot 14 + 80 \cdot 8 = 1900 = \tilde{z}$. ✓
4. Reduced costs $\bar{c}_j = c_j - \sum_i a_{ij}\, \tilde{\pi}_i$: $\quad \bar{c}_1 = 30 - (14 + 16) = 0, \quad \bar{c}_2 = 50 - (42 + 8) = 0$ (zero: both variables are nonzero — beware of degeneracy: there can also be variables at zero with $\bar{c} = 0$).

Sensitivity on the 2×2 example: shadow prices and reduced costs

Shadow prices

Increase b_1 by one unit at a price of 12? Worth it $(14 > 12)$. With $b_1 = 100$: $2040 = 1900 + 10 \cdot 14$ — exact as long as b_1 stays in the validity range $[b_1^{\min}, b_1^{\max}] = [40, 240]$. A non-binding constraint always has zero shadow price.

Reduced costs: is a third variable worth it?

Coefficient $c_3 = 20$, consumptions $a_{13} = a_{23} = 1$:

- at the shadow prices the resources are worth $1 \cdot 14 + 1 \cdot 8 = 22 > 20$;
- solving: $\tilde{x}_3 = 0$ with $\bar{c}_3 = 20 - 22 = -2$ and $c_3^{\max} = 22 \Rightarrow$ worth it only from $c_3 = 22$ upwards;
- cross-check with $c_3 = 23$: the solution changes to $(0, 5, 75)$, value 1975;
- validity ranges of the coefficients: c_1 may vary in $[16.7,\ 100]$ without changing the solution.

All cases in a single LP: checking the conditions

- Solving gives: $(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3) = (60, 40, 0)$, $\tilde{z} = 620$; $(\tilde{\pi}_1, \tilde{\pi}_2, \tilde{\pi}_3) = (0,\ 6,\ -1)$; $(\bar{c}_1, \bar{c}_2, \bar{c}_3) = (0,\ 0,\ -3)$; coefficient ranges $(-\infty, 8]$, $[5, 17]$, $(-\infty, -6]$.
- Checks: $\tilde{\pi}_1 = 0$ (slack constraint: $100 > 30$); $\tilde{\pi}_2 = 6$ (equality, free dual); $\tilde{\pi}_3 = -1 \leq 0$ ($\leq$ in a minimization); strong duality $600 + 20 = 620$. ✓
- x_3 at its upper bound (zero): $\bar{c}_3 = -9 + 6 = -3$, threshold $c_3^{\max} = -6$; perturbations: $b_2{+}1 \to 626$, $b_3{+}1 \to 619$, $x_3 = -1 \to 623$. ✓

Theory: nonlinear optimization

A 2×2 QP, worked out in full

Separable quadratic objective, one threshold constraint on the sum:

$$\begin{array}{rrcrcl} \min & x_1^2 & + & 2\,x_2^2 & & \\ \text{subject to} & x_1 & + & x_2 & \geq & 6, \\ & x_1, & & x_2 & \geq & 0. \end{array}$$

- $\boldsymbol{Q} = \operatorname{diag}(2, 4) \succ 0$: convex $\Rightarrow$ certified global optimum.
- Solving gives: $(\tilde{x}_1, \tilde{x}_2) = (4, 2)$, $\tilde{f} = 24$, $\lambda = 8$.
- Check: binding constraint $(4 + 2 = 6)$ and equal derivatives $2\tilde{x}_1 = 8 = 4\tilde{x}_2 =$ shadow price of the threshold. ✓

KKT conditions

For $\min f(x_1,\ldots,x_n)$ subject to $g_i(x_1,\ldots,x_n) \leq 0$, $h_j(x_1,\ldots,x_n) = 0$, with Lagrangian $L = f + \sum_i \lambda_i g_i + \sum_j \nu_j h_j$ ($\lambda_i \geq 0$), at a regular optimum $(\tilde{x}_1,\ldots,\tilde{x}_n)$:

$$\frac{\partial L}{\partial x_k}(\tilde{x}_1,\ldots,\tilde{x}_n) = 0, \quad \forall k \in \{1,2,\ldots,n\} \qquad \text{(stationarity)}$$

$$\lambda_i\, g_i(\tilde{x}_1,\ldots,\tilde{x}_n) = 0, \quad \forall i \in \{1,2,\ldots,m\} \qquad \text{(complementarity)}$$

If the problem is convex, they are also sufficient.

Worked example: KKT on the example QP

$\min x_1^2 + 2x_2^2$ subject to $x_1 + x_2 \geq 6$. Stationarity: $2x_1 = \lambda$, $4x_2 = \lambda$; the constraint is binding, so $\lambda = 8$, $(\tilde{x}_1, \tilde{x}_2) = (4,2)$, $\tilde{f} = 24$; with $6+\varepsilon$: $\tilde{f} = 24 + 8\varepsilon + \frac{2}{3}\varepsilon^2$.

- In the linear case the KKT conditions are the complementary slackness of (P)–(D): multipliers = duals π_i.
- No duals available: estimate them by perturbation.

The sensitivity analysis protocol (in every lab session)

1. **Base scenario**: solve, verify, binding constraints.
2. **One-at-a-time**: one key parameter over a grid.
3. **Shadow prices and reduced costs**: dual vs perturbed re-optimization; $\bar{c}_j$ of the variables at zero $=$ break-even threshold.
4. **Scenarios**: pessimistic, central, optimiztic.
5. **Trade-off**: a frontier (cost-service, risk-return, cost-CO_2).
6. **Stability**: data $\pm 5\%$.

Implementation: linear models

Building a model in five steps

```
m = gp.Model("production")                    # 1. container
x = m.addVars(products, months, name="x")     # 2. variables (>= 0)
m.addConstrs((x.sum("*", t) <= b[t]           # 3. constraints
              for t in months), name="cap")
m.setObjective(..., GRB.MINIMIZE)             # 4. objective
m.write("model.lp"); m.optimize()             # 5. debug + solution
```

- Free variables (SVM: b; CVaR: η): `m.addVar(lb=-GRB.INFINITY)` — forgetting it is the most common error!
- Sums: `gp.quicksum(...)` or `x.sum(i, "*")`.

Reading the solution

```
if m.Status == GRB.OPTIMAL:      # ALWAYS before reading numbers
    print(m.ObjVal, x[i,t].X)
elif m.Status == GRB.INFEASIBLE:
    m.computeIIS(); m.write("conflict.ilp")
```

`Pi`	shadow price of the constraint (LP)
`Slack`	slack: $0 \Rightarrow$ binding constraint
`RC`	reduced cost of a variable at zero
`SARHSLow/Up`	validity range of the shadow price
`SAObjLow/Up`	validity range of the reduced cost

- Minimization problems, $\leq$ constraint: `Pi` ≤ 0 (Gurobi convention).
- RHS with constants on the left: Gurobi moves them into the right-hand side (pitfall in sensitivity analysis — see EV charging).

Interpretation checklist

1. Status `OPTIMAL`? If not: diagnose (IIS, `.lp` file).
2. Is the order of magnitude of the optimum sensible?
3. Binding constraints $=$ the expected ones?
4. Shadow prices: which resource to expand first, and up to what price?
5. Stability: data $\pm 5\%$.
6. Managerial recommendation in three lines.

Implementation: nonlinear models

One solver only: general NLPs in Gurobi too

- Nonlinear functions as function constraints on auxiliary variables: `addGenConstrLog/Exp/Pow` with `FuncNonlinear=1`.
- Bilinear terms and ratios with an auxiliary variable and `NonConvex=2` (e.g. $w\,(\mu - \lambda) = 1$).
- Certified global optimum, same interpretation checklist.
- For marginal analyses: `MIPGap`, `FeasibilityTol`, `OptimalityTol` at 10^{-9}; reformulate for scaling (e.g. $q \cdot p^{\varepsilon} \leq A$, not $q \leq A\,p^{-\varepsilon}$).

Multi-period production and inventory (LP/QP)

Production and inventory — the problem

- **Decisions**: production x_{it} and inventory s_{it} (nonnegative); **objective**: minimum cost.
- **Data**: demand d_{it}, costs c_{it}, h_{it}, hours per unit a_i, hours available b_t, initial inventory s_{i0}.

Model (LP)

$$
\begin{aligned}
\min \quad & \sum_{i=1}^{n} \sum_{t=1}^{m} (c_{it}\, x_{it} + h_{it}\, s_{it}) \\
\text{subject to} \quad & s_{i,t-1} + x_{it} = d_{it} + s_{it}, && \forall i \in \{1,2,\ldots,n\},\ \forall t \in \{1,2,\ldots,m\}, \\
& \sum_{i=1}^{n} a_i\, x_{it} \le b_t, && \forall t \in \{1,2,\ldots,m\}, \\
& x_{it},\ s_{it} \ge 0, && \forall i \in \{1,2,\ldots,n\},\ \forall t \in \{1,2,\ldots,m\}.
\end{aligned}
$$

Production and inventory — understanding the constraints

- **Objective**: the model freely trades early production (more inventory) against production at the capacity limit.
- **Balance**: inventory accounting; it is the constraint that links the months: without it the problem splits into $|T|$ independent problems.
- **Capacity**: the shared scarce resource; its dual says how much one extra hour is worth.

Worked example by hand (1 product, 2 months)

$d = (60, 100)$, capacity 80/month, $c = 10$, $h = 1$: the only feasible plan is $x = (80, 80)$, $s_1 = 20$, cost 1.620 €. $+1$ hour in month 2 $\Rightarrow$ cost -1 € (one month of holding less): dual of month 2 $= -1$, month 1 $= 0$.

Production and inventory — case study

3 products $\times$ 6 months, demand peak in months 4–5, 420–460 hours/month.

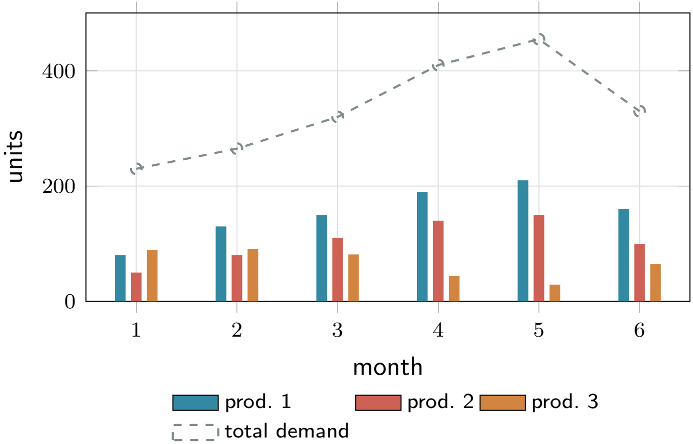

- Optimal cost: **32.889 €**
- Product 3 (2.1 hours/unit) is produced early (*pre-build*) and stocked (up to 107 units) to free up the peak months.
- Zero inventory in the last month: the model does not waste.

Production and inventory — the shadow prices tell a story

month	1	2	3	4	5	6
hours used	330/420	420/420	460/460	460/460	460/460	420/420
shadow price (€/h)	0	-0.76	-1.52	-2.29	-3.05	-3.81

- The duals grow by $0.762 = h_3/a_3$ per month: one extra hour in month t saves one month of holding product 3.
- Check by perturbation: $+1$ hour in month 6 $\Rightarrow$ cost -3.810 € (it matches the dual).
- When the solver and economic reasoning give the same number, the model is probably right.

Production and inventory — smoothing and the value of capacity

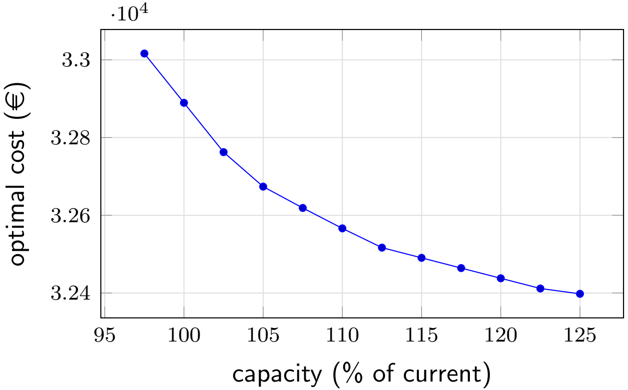

- **Convex and decreasing** curve: the first extra hours are worth a lot.
- Below 97%: infeasible — demand can no longer be met.
- QP variant with $\gamma \sum_{t>1} (X_t - X_{t-1})^2$: an almost flat profile (max change from 81 down to 1 unit) for $+114$ € $(+0.35\%)$.

Take-away

Overtime in months 5–6 pays off up to 3–3.8 €/hour below the internal price; the levelled plan is almost free.

Supply chain with congestion and CO_2 (LP/NLP)

Supply chain — the problem

In words. Decide how many units to ship on each leg of the network. Objective: minimum transport cost. Constraints: flow conservation at every node and capacity of the legs.

Model (minimum cost flow)

$$\min \sum_{(i,j)\in A} c_{ij}x_{ij} \text{ subject to } \sum_{j:(i,j)\in A} x_{ij} - \sum_{j:(j,i)\in A} x_{ji} = b_i,\ \forall i \in N,\ \ 0 \le x_{ij} \le u_{ij}$$

Objective variants: quadratic congestion $\sum(c_{ij}x_{ij} + \alpha c_{ij}x_{ij}^2/u_{ij})$ (convex); internal CO_2 price: $\sum(c_{ij} + \tau e_{ij})x_{ij}$; minimax of the utilization.

- Worked example (2 routes): the LP saturates the cheap route; congestion equalises the marginal costs and leaves a safety margin.

Supply chain — case study

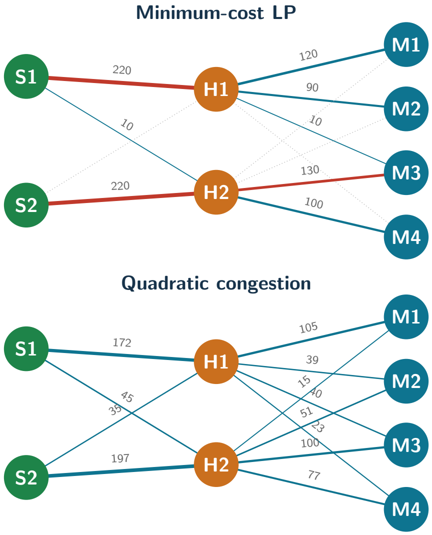

2 plants → 2 hubs → 4 markets; cheap "road" legs (polluting), expensive "rail" ones (clean).

- LP: cost 3.385 €, 3 saturated arcs (in red), emissions 2.730 kg.
- Congestion: cost 3.703 €, no arc above 90%.
- Shadow prices of demand = marginal costs of serving the markets: M1 8.00 — M4 10.50 €/unit.

Supply chain — the price of CO_2

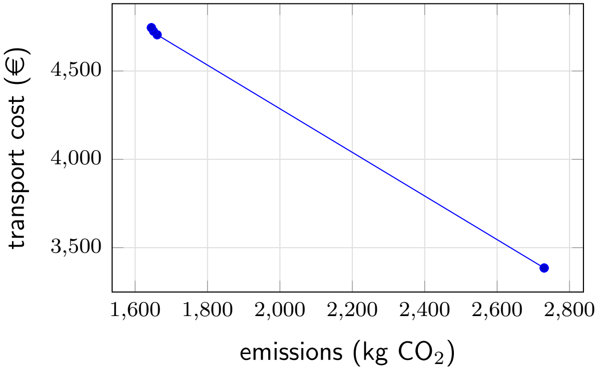

- $\tau \leq 1$: "road" configuration: 3.385 € and 2.730 kg.
- $\tau \geq 1.5$: "rail" configuration: 4.705 € and 1.661 kg.
- Exact threshold, by bisection: $\tau^* = 1.25$ €/kg.
- LP solutions sit at the vertices: the frontier moves in jumps; below the threshold the internal price changes nothing.

Warning

Pure minimax $(\min z)$ is degenerate: every routing with utilization $\leq z^*$ is optimal, however expensive. Always combine it with the cost.

Markowitz portfolio (QP)

Markowitz — the problem

In words. Split the capital across n assets. Objective: minimum variance. Constraints: weights summing to 1, minimum expected return, per-asset limits.

Model (convex QP: $\boldsymbol{Q} \succeq 0$ always)

$$\min \sum_{i=1}^{n}\sum_{j=1}^{n} q_{ij}\, x_i x_j \quad \text{subject to} \quad \sum_{i=1}^{n} \mu_i x_i \geq \bar{r}, \quad \sum_{i=1}^{n} x_i = 1, \quad \ell_i \leq x_i \leq u_i$$

Worked example (2 uncorrelated assets)

$\sigma_1 = 20\%$, $\sigma_2 = 30\%$: $x_1 = 9/13 = 69.2\%$, portfolio volatility 16.6% — less than either asset: diversification in its purest form.

Markowitz — the efficient frontier

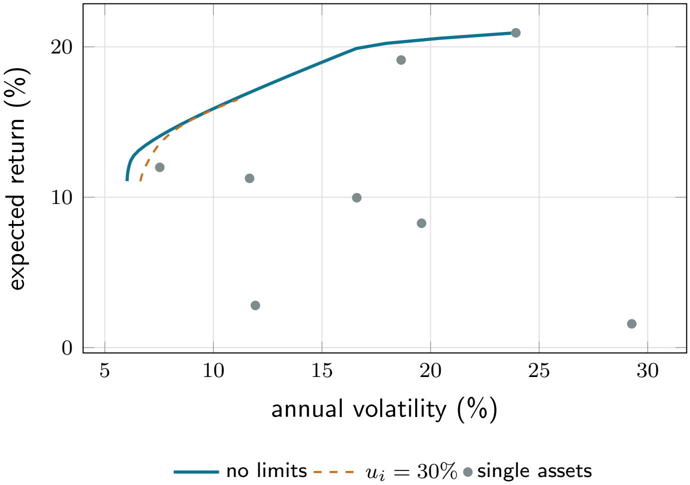

8 ETFs, μ and $\boldsymbol{Q}$ estimated from 60 months.

- Minimum variance: return 11%, volatility **6%** $<$ 7.5% of the best single asset.
- Equally weighted $1/n$: dominated (10.7%, 10.1%).
- Cap $u_i = 30\%$: it cuts the top — the cost of the mandates *is visible*.

Markowitz — sensitivity and fragility

- The return constraint is non-binding up to $\bar{r} = 11.06\%$ (minimum variance already yields a lot): asking for "at least 8%" costs nothing.
- Beyond that: every point of return is paid for in variance (multiplier ≈ 0.022 at $\bar{r} = 12\%$).
- Composition along the frontier: from the defensive mix (UTL, CON, SAN) to a portfolio concentrated in the aggressive assets (MAT, ENE).

The most important practical lesson

Return estimates are far more unstable than covariance estimates (standard error $\approx 2.6\%$ per year over 60 months): optimized portfolios chase estimation errors. Remedies: u_i limits, shrinkage, minimum variance.

Pricing and revenue management (NLP)

Pricing — the problem

In words. Decide price p and quantity q. Objective: maximum profit $(p-c)q$. Constraints: $q \le d(p)$ and $q \le k$ (capacity).

Model and demand curves

$$\max\ p\,q - c\,q \ \text{ subject to }\ q \le d(p),\ q \le k$$

linear $d(p) = a - bp$; constant elasticity $d(p) = \theta p^{-\varepsilon}$; logistic $d(p) = m/(1 + e^{-\alpha+\beta p})$.

- The nonconvexity lies in the product $p \cdot q$ (bilinear); with linear demand the reduced profit is a concave parabola; Gurobi still solves it to global optimality (`NonConvex=2`).

Pricing — worked example and case study

Concert: $d(p) = 1200 - 5p$, $c = 20$ €, $k = 400$ seats.

1. Without capacity: $p^\circ = (a/b + c)/2 = 130$ €, $q^\circ = 550$.
2. Capacity bites: $p^* = (a - k)/b =$ **160** €, profit 56.000 €.
3. Value of one extra seat: $(p^* - c) + k\,dp^*/dk = 140 - 80 =$ **60** € — not the full margin: filling the seat means lowering the price for everyone. (Gurobi, by perturbation: 59.80.)

The optimum at the kink

With constant elasticity the unconstrained optimum would be 36.67 €, but there demand explodes: the true optimum is at the kink $d(p) = k$, i.e. $p^* = (\theta/k)^{1/\varepsilon} = 79.11$ € — no derivative vanishes.

Pricing — profit and seating capacity

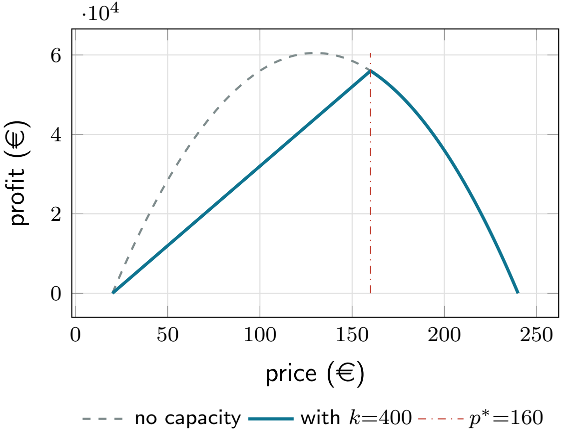

- The capacity constraint cuts the parabola and moves the optimum from 130 to 160 €.
- The value-of-capacity curve flattens at $q^\circ = 550$: beyond that, seats are worth nothing.
- Multi-product (stalls/balcony with substitution): prices to be set jointly; profit 85.149 €.
- Expansions: the balcony pays ($+3.838$ € for 50 seats), the stalls do not ($+3.287$ €).

Advertising budget (convex NLP)

Advertising budget — the problem and the KKT condition

In words. Split the budget b across channels with concave response $r_i(x_i) = a_i \log(1 + k_i x_i)$; constraints: budget and caps u_i.

$$\max \sum_{i=1}^{n} r_i(x_i) \quad \text{subject to} \quad \sum_{i=1}^{n} x_i \leq b, \quad 0 \leq x_i \leq u_i$$

KKT = economics

At the optimum $r_i'(x_i^*) = \lambda$ for every interior channel: the last euro returns the same in every active channel. Channels at their cap: marginal $> \lambda$; channels at zero: initial marginal $< \lambda$.

Worked example ($b = 50$): $x = (35.6;\, 14.4)$, $\lambda = 2.46$.

Advertising budget — numerical check and value curve

	social	search	TV	influencer
spend (thousand €)	24.3	34.8	22.9	18.0
marginal return	8.2693	8.2693	8.2693	8.2693

- Check: $+1.000$ € of budget $\Rightarrow$ response $+8.244 \approx \lambda$.
- TV gets less than social despite the higher cap: what matters is the saturation speed k_i, not the size of the channel.
- λ falls from 19 ($b = 20$) to 0 ($b = 300$, all at their caps): the λ curve is the argument for negotiating the budget with the CFO.

Continuous location (convex NLP)

Location — three objectives, three answers

In words. Decide the coordinates (x, y) of the facility; the objective — average, maximum or quadratic — is the real managerial choice.

$$\text{Weber:} \quad \min \; \sum_{i=1}^{n} w_i \sqrt{(x-a_i)^2 + (y-b_i)^2}$$

$$\text{Minimax:} \quad \min \; z \quad \text{subject to} \quad \sqrt{(x-a_i)^2 + (y-b_i)^2} \le z, \;\; \forall i \in \{1, 2, \ldots, n\}$$

Worked example in 1D (weights 1, 1, 3 at km 0, 4, 10)

Weber $=$ weighted median $= 10$; centroid $= 6.8$; minimax $= 5$. Three reasonable objectives, three different answers.

- Quadratic $\Rightarrow$ weighted centroid (closed form); Weber: not differentiable at the demand points; minimax: centre of the smallest enclosing circle, it ignores the weights.

Location — case study and equity

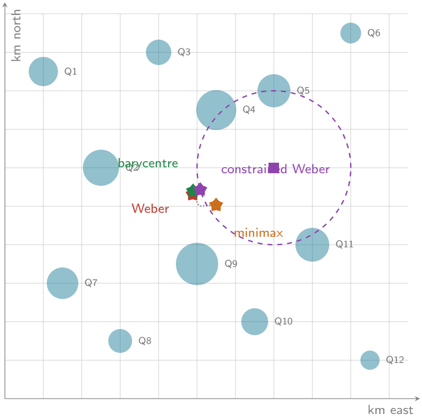

12 districts; constraint: within 2 km of the substation at $(7, 6)$.

- Weber $(4.89; 5.31)$: average 3.344 km, max 6.314.
- Minimax $(5.50; 5.03)$: average 3.391, max **5.680**.
- Substation constraint: $+0.2\%$ — almost free.
- The efficiency-equity frontier (ε-constraint) is almost flat: -0.63 km for the farthest district costs 47 m of average.

EV charging (LP/QP)

EV charging — the problem

In words. How much power x_{vt} per vehicle and hour, at minimum cost; availability $a_{vt} \in \{0.1\}$ is data, not a decision.

Model (LP: no binary variables!)

$$
\begin{aligned}
\min \quad & \sum_{v=1}^{n} \sum_{t=1}^{m} \pi_t \, \Delta t \, x_{vt} \\
\text{subject to} \quad & \eta \sum_{t=1}^{m} \Delta t \, x_{vt} \geq e_v, && \forall v \in \{1, 2, \ldots, n\}, \\
& \sum_{v=1}^{n} x_{vt} + b_t \leq k, && \forall t \in \{1, 2, \ldots, m\}, \\
& 0 \leq x_{vt} \leq a_{vt} \, \bar{p}_v, && \forall v \in \{1, 2, \ldots, n\}, \ \forall t \in \{1, 2, \ldots, m\}.
\end{aligned}
$$

EV charging — minimum cost vs peak shaving

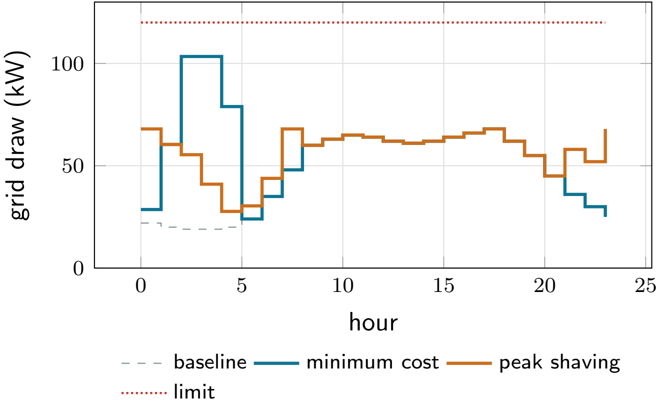

- Minimum cost: 20.36 €/night, peak **103 kW**.
- Peak shaving: peak 68 kW; with the cost$+\rho$ peak compromise the same peak costs 22.15 € (pure minimax: 27.79!).
- Duals of the energy needs $= 0.08/\eta$ or $0.09/\eta$: the price of each vehicle's marginal hour.
- Minimum feasible capacity: 68 kW (bisection).

A real pitfall

In `quicksum(x) + base[t] <= C` the stored RHS is `C - base[t]`.

Queues and service capacity (convex NLP)

M/M/1 queues — the problem

In words. How much service capacity to buy; objective: capacity cost + value of the customers' time; constraint: stability $\mu > \lambda$.

Model and closed-form solution

$$\min_{\mu>\lambda}\ c\,\mu + h\,\frac{\lambda}{\mu-\lambda} \qquad \Longrightarrow \qquad \mu^* = \lambda + \sqrt{h\lambda/c}$$

"capacity = demand + buffer"; the buffer grows like $\sqrt{\lambda}$ (economies of scale, pooling).

Worked example ($\lambda = 8$, $c = 2$, $h = 4$): $\mu^* = 12$, $\rho = 67\%$, cost 32 €/h; with $\mu = 9$ ("more efficient", $\rho = 89\%$): 50 €/h.

Queues — the utilization wall and the price of a promise

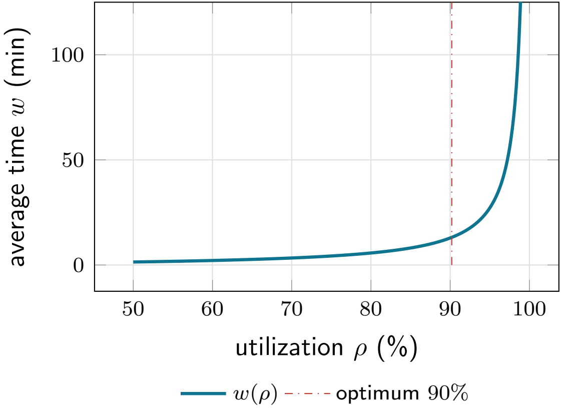

Case study: $\lambda = 42$, $c = 3$, $h = 1.5 \Rightarrow \mu^* = 46.6$, $\rho = 90\%$, $w = 13$ min.

- Between $\rho = 90\%$ and 99% the wait multiplies by **ten**: a mathematical conflict, not an organizational one.
- Promise $w \leq w_{\max}$: from 12 to 9 min costs 1.85 €/h; from 3 to 2 min 28.95 €/h ($dC/dw_{\max} = -c/w_{\max}^2 + h\lambda$).
- λ uncertain in $[36, 48]$: the nominal optimum breaks down ($\mu^* < 48$); the robust one ($\mu = 52.9$) costs $+13\%$.

The newsvendor (stochastic LP)

Newsvendor — deciding before knowing

In words. A single quantity q, chosen *before* observing the demand D; objective: minimum expected cost of overage (c_o) and underage (c_u).

Model and the quantile rule

$$\min_{q \geq 0} \; c_o \mathbb{E}[(q-D)^+] + c_u \mathbb{E}[(D-q)^+] \qquad \Longrightarrow \qquad F(q^*) = \frac{c_u}{c_u + c_o}$$

Worked example ($D \in \{80, 100, 120\}$, $c_u = 9$, $c_o = 4$)

$C(80) = 180$, $C(100) = 86.7$, $C(120) =$ **80**: one orders the maximum, not the mean — $F(100) = 0.67 < \alpha^* = 0.692$.

Continuous case ($\mu = 100$, $\sigma = 20$, $c_u/c_o = 9/4$): $q^* = F^{-1}(0.6923) = 110 \neq 100$.

Newsvendor — scenario LP and value of the stochastic solution

The positive-part trick (no binaries)

$$\min\ c_o \sum_{s=1}^{k} \pi_s o_s + c_u \sum_{s=1}^{k} \pi_s u_s \quad \text{subject to}\ \ o_s \geq q - d_s,\ \ u_s \geq d_s - q,\ \ q, o_s, u_s \geq 0$$

At the optimum $o_s = (q - d_s)^+$ and $u_s = (d_s - q)^+$: the objective squeezes them.

- The q of the LP = the empirical quantile (108.6 with 600 scenarios): a theorem.
- **VSS** $= 11.17$ €/cycle (-11%): the gain from modelling uncertainty instead of ordering "the mean".
- Stability: with 30 scenarios q moves by ± 3.5 units across replications — the scenarios are a sample.
- Service level: 95% no-stockout costs $+45.59$ €; the fill rate at q^* is already 96% — read the SLA carefully.

Newsvendor — risk aversion and multi-product

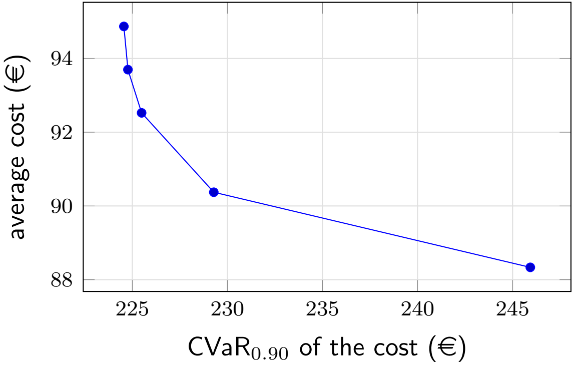

- Mean-CVaR: from $\lambda = 0$ to 1 one orders more (108.6 $\rightarrow$ 116.2): $+6.5$ € of average buy -21.4 € of tail.
- Steep frontier at the start: much protection almost for free.
- Multi-product with a 1.200 € budget and correlated demands ($\rho = 0.7$): budget dual -0.55 (a 55% return!); with $\rho = 0$ the expected cost falls by 18%: correlation is a cost.

VaR and CVaR (LP)

VaR and CVaR — measuring the tail

L random loss, $\alpha \in (0.1)$:

$$\mathrm{VaR}_\alpha(L) = \inf\{\eta : \mathbb{P}(L \leq \eta) \geq \alpha\},$$
$$\mathrm{CVaR}_\alpha(L) = \text{average of the worst tail of mass } 1-\alpha.$$

- CVaR is convex and subadditive (it rewards diversification); VaR is not.
- Worked example (losses $\{2.4, 5, 7, 12, 20\}$, $\alpha = 0.8$): VaR $= 12$; CVaR $= 12 + \frac{1}{0.2} \cdot \frac{1}{6}(20-12) =$ **18.67**; average $= 8.33$.
- Three numbers, three stories: the average reassures, the VaR gives a threshold, the CVaR says how much it hurts when things go badly.

The Rockafellar–Uryasev linear formulation

Scenarios $s \in \{1, 2, \ldots, k\}$ with probabilities π_s; linear loss $\ell_s(\boldsymbol{x})$; threshold η (a *free* variable) and excesses $\xi_s \geq 0$.

LP model

$$\begin{aligned} \min \quad & \eta + \frac{1}{1-\alpha} \sum_{s=1}^{k} \pi_s \, \xi_s \\ \text{subject to} \quad & \xi_s \geq \ell_s(\boldsymbol{x}) - \eta, && \forall s \in \{1, 2, \ldots, k\}, \\ & \xi_s \geq 0, && \forall s \in \{1, 2, \ldots, k\}. \end{aligned}$$

- At the optimum η^* is a VaR and the objective is the CVaR: a single LP, both measures.
- Same positive-part trick as the newsvendor: only the scenarios in the tail contribute.

VaR and CVaR — mean-CVaR portfolio vs Markowitz

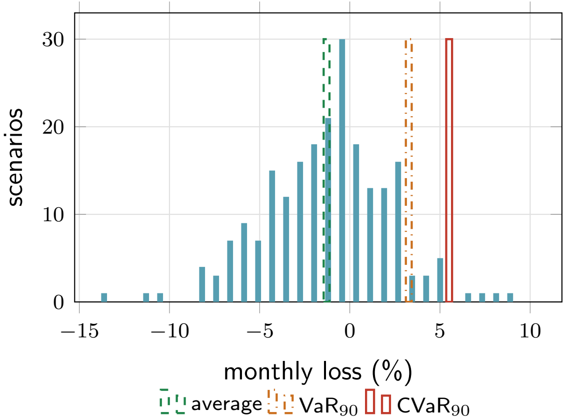

220 scenarios with fat tails (Student t), required return 8%.

- Same average loss, but CVaR_{90}: mean-CVaR 4.89% vs Markowitz 5.31% (-8%).
- Variance penalizes above and below the average; CVaR looks only where it hurts.
- With $\alpha = 0.99$ and 220 scenarios the tail has 2–3 points: unstable estimates — scenarios beyond the quantile are needed.

VaR and CVaR — two-stage supply chain: the cost of resilience

First stage: capacity x_a reserved from the suppliers (F1 cheap but fragile, F2 expensive but reliable). Second stage (per scenario): flows and shortages.

λ	reserved F1	reserved F2	average cost	CVaR_{90}
0 (neutral)	161.5	134.2	1.724	2.902
0.5	92.7	212.7	1.803	2.217
1 (protected)	68.5	236.9	1.906	2.139

Take-away

As risk aversion grows, capacity migrates towards the reliable supplier: $+79$ € of average cost buy -685 € of CVaR — the number needed in discussions on dual sourcing and reshoring.

Arbitrage and pricing (LP)

Arbitrage — the problem

- One-period market: today prices s_i^0; tomorrow one of m states of the world, asset i pays s_{ij}^1; asset 0 is risk free ($s_0^0 = 1$, it pays $R = 1 + r$ everywhere).
- Arbitrage: cash in today with no risk tomorrow (type A), or free today and never at a loss, with a chance of a gain (type B).
- Detecting it is an LP; duality is the theory of pricing.

The arbitrage detection LP

Free positions x_i (buying or short selling)

$$\begin{aligned} \min \quad & \sum_{i=0}^{n} s_i^0 \, x_i \\ \text{subject to} \quad & \sum_{i=0}^{n} s_{ij}^1 \, x_i \geq 0, && \forall j \in \{1, 2, \ldots, m\}, \\ & x_i \gtreqless 0, && \forall i \in \{0, 1, \ldots, n\}. \end{aligned}$$

- The zero portfolio is feasible: the optimum is always ≤ 0.
- If < 0: type A and the model is unbounded (homogeneous constraints) $\Rightarrow$ normalization $\sum_i s_i^0 x_i \geq -1$ ("cash in at most 1 today").

Duality is pricing

Every state gives a $p_j \geq 0$; every free position gives an equality dual constraint:

$$\sum_{j=1}^{m} s^1_{ij}\, p_j = s^0_i \quad \forall i \in \{0, 1, \dots, n\}, \qquad p_j \geq 0.$$

From asset 0: $\sum_j p_j = 1/R \Rightarrow q_j = R\, p_j$ are risk-neutral probabilities and

$$s^0_i = \frac{1}{R} \sum_{j=1}^{m} q_j\, s^1_{ij} :$$

today's price is the discounted expected value of the payoffs under q.

The fundamental theorem of asset pricing

Theorem

No arbitrage (neither type A nor type B) $\iff$ there exists a risk-neutral measure with $q_j > 0$ for every state $j \in \{1, 2, \dots, m\}$.

- No type A $\Rightarrow$ optimum 0 $\Rightarrow$ by strong duality the dual is feasible: a measure exists.
- No type B $\Rightarrow$ all the state constraints are binding $\Rightarrow$ by complementary slackness there is a measure with $q > 0$.
- Conversely, with $q > 0$ every riskless strategy costs ≥ 0.

Arbitrage — case study

Three states, $R = 1.04$; payoff of asset 1 $(10, 15, 13)$, asset 2 $(30, 15, 25)$.

- **Prices (6, 20)**: UNBOUNDED; normalized: optimum -1. By hand: $(-27, 1, 1)$ cashes in 1 today, payoff $(11.92;\ 1.92;\ 9.92) \geq 0$.
- **Prices (13, 18.69)**: optimum 0; from the duals of the state constraints $q = (0.296;\ 0.704;\ 0)$ — and the s_i^0 are reproduced exactly.
- **Pricing a call** (asset 1, strike 12, payoff $(0.3, 1)$): complete market $\Rightarrow$ unique price 2.0308; incomplete (assets 0 and 1 only) $\Rightarrow$ range $[1.4615,\ 2.0308]$.

Arbitrage — the range of consistent prices

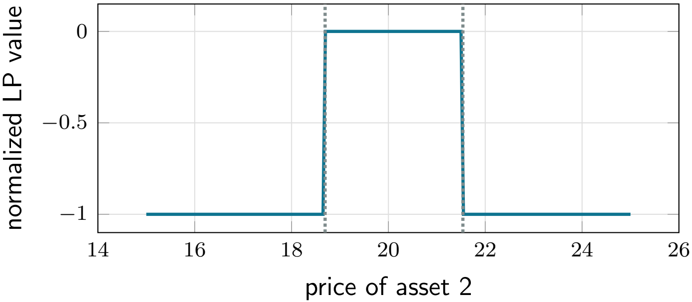

- Two LPs (min/max over the consistent measures): $s_2^0 \in [18.69, 21.54]$; inside, value 0; outside, -1: arbitrage in one direction or the other.

Support Vector Machine (QP)

SVM — the classifier as a QP

In words. Find the hyperplane that separates two classes maximizing the margin; the labels $y_i \in \{-1, +1\}$ are data, the variables ($\boldsymbol{w}$, b, slacks ξ_i) are all continuous.

Soft margin (primal)

$$\min_{\boldsymbol{w},b,\boldsymbol{\xi}} \ \tfrac{1}{2}\|\boldsymbol{w}\|^2 + C\sum_{i=1}^{n}\xi_i \quad \text{subject to} \quad y_i\Big(\sum_{j=1}^{p} w_j x_{ij} + b\Big) \ge 1 - \xi_i, \quad \xi_i \ge 0$$

margin $= 2/\|\boldsymbol{w}\|$; $C =$ price of the violations.

Worked example (2 points): $\boldsymbol{w} = (0.5; 0.5)$, $b = -1$, margin $2\sqrt{2} =$ distance between the points; both are support vectors.

SVM — case study: credit risk

90 customers, two indicators; as C varies:

	$C = 0.05$	$C = 1$	$C = 20$
margin	2.42	1.28	0.36
training errors	2	2	0
points inside the margin	22	7	3

- **Dual**: $\max \sum_i \alpha_i - \frac{1}{2} \sum_{i,j} \alpha_i \alpha_j y_i y_j \boldsymbol{x}_i' \boldsymbol{x}_j$, $0 \le \alpha_i \le C$, $\sum_i \alpha_i y_i = 0$: same $\boldsymbol{w}$, b and value (strong duality, 4.7868).
- Complementarity: $\alpha_i = 0$ (83 points, irrelevant); $0 < \alpha_i < C$ (2, on the margin: they give b); $\alpha_i = C$ (5, inside the margin).
- Imbalanced classes: cost $10\times$ on the defaulters $\Rightarrow$ recall 100% at the price of precision 97%.

SVM — kernels and SVR

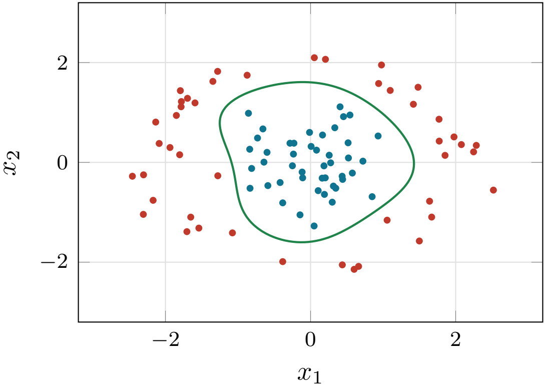

- The dual depends on the data only through $\boldsymbol{x}_i'\boldsymbol{x}_j$: replacing it with $K(\boldsymbol{x},\boldsymbol{z}) = e^{-\gamma\|\boldsymbol{x}-\boldsymbol{z}\|^2}$ gives curved boundaries while staying in a convex QP.
- Small γ: smooth boundary; $\gamma = 5$: 61 support vectors out of 90, the model memorises (overfitting). C and γ by cross-validation.
- **SVR**: a $\pm\varepsilon$ tube with no penalty; on the price-demand data it recovers $-10.86p + 220.3$ (true: $-11p + 220$) using only the 14 points outside the tube.

Robust and quantile regression (LP)

Regression — estimating the parameters with an LP

In words. Choose the coefficients of a linear function that predicts demand; the objective pays τ for the deviations above and $1-\tau$ for those below (u_i and v_i, the two parts of the deviation).

Quantile regression

$$\min_{\boldsymbol{w},b,\boldsymbol{u},\boldsymbol{v}} \sum_{i=1}^{n} \left[\tau u_i + (1-\tau)v_i\right] \quad \text{subject to} \quad \sum_{j=1}^{p} w_j x_{ij} + b + u_i - v_i = y_i, \quad u_i, v_i \geq 0$$

w_j and b free; $\tau = 1/2$ gives the median (absolute deviations), a large τ a prudent prediction on the high side.

With no features ($p = 0$) the optimum is the empirical quantile: with $88, 96, 104, 112, 120, 132, 148$ and $\tau = 9/13$ we get $\tilde{b} = 120$ and $\tilde{z} = 680/13 = 52.31$ — it is the newsvendor quantile rule.

Regression — robustness and support points

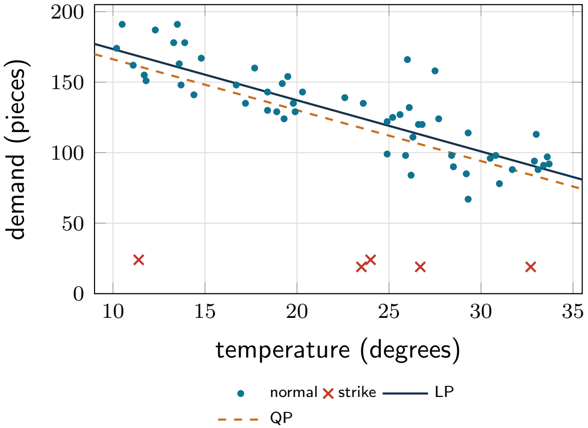

- Five strike days out of 60: least squares pay the deviation squared and the intercept drops by 11 pieces; the LP stays where the normal data are ($209.77 - 3.63\,t$ against $213.58 - 3.72\,t$ of the QP cleaned by hand).
- **Duals**: $-(1-\tau) \leq \tilde{\pi}_i \leq \tau$; above the estimate $\tilde{\pi}_i = \tau$, below $-(1-\tau)$, and only the $p+1$ support points have an interior dual — the support vectors of regression.
- $\sum_i \tilde{\pi}_i = 0$ (dual of the intercept) *is* the quantile property: 29 above, 29 below, 2 interpolated.
- Typical day: median 126.3, $9/13$ quantile 132.1 against 148 for the unconditional quantile.

Lab organization

The four labs and the deliverable

Lab 1	Production and inventory	LP, duals, checking the shadow prices
Lab 2	Markowitz	QP, frontier, fragility of the estimates
Lab 3	Pricing *or* budget	NLP, concavity, numerical KKT
Lab 4	Project of your choice	managerial presentation

Report (max 8 pages): problem and assumptions → model (constraints explained) → data → results → sensitivity (full protocol) → managerial recommendation in ten lines with no formulas.

Grading: formulation 30% · implementation and verification 25% · sensitivity 25% · interpretation 20%.

The most common mistakes (checklist)

1. Reading `.X` or `.Pi` without checking `m.Status`.
2. Forgetting `lb=-GRB.INFINITY` on the free variables (b, η).
3. Using a shadow price outside its validity range.
4. Getting the sign of the duals wrong in minimization problems.
5. Updating the RHS of a constraint with constants on the left.
6. Optimizing a single objective when there are two (pure minimax).
7. Choosing the hyperparameters by looking at the test set.
8. Six decimal places from estimates that wobble at the second.

Materials

- **Full lecture notes**: theory, examples worked out by hand, case studies, code, sensitivity, exercises (with solutions for instructors).
- **Python scripts**: one per chapter; `python/run_all.py` regenerates data, results and figures.
- **Solver guide**: `GUIDA_GUROBI.md`.
- **Data** in CSV with fixed seeds: everything is reproducible.

Enjoy the lab!