SEN605 · WEEK 8

Modern Portfolio Theory and the Efficient Frontier

Markowitz mean-variance optimization · Sharpe ratio · Monte Carlo portfolio clouds

Harry Markowitz (1952) · William Sharpe (1964) · Fischer Black (1972)

Assoc. Prof. Dr. Onur Polat

Hacettepe University · Graduate School of Informatics · Information Systems

Today's Agenda

6 blocks · 2 hours · Python + pandas + matplotlib hands-on

01

Why Diversification Works

The one free lunch in finance · 15 min

02

Return and Risk of a Portfolio

Expected return, variance, covariance · 15 min

03

The Efficient Frontier

Markowitz mean-variance framework · 20 min

04

The Sharpe Ratio & Tangent Portfolio

Capital Market Line · 15 min

05

BIST + US Live Lab

Build the frontier from real data · 30 min

06

Practical Limits & Extensions

Shrinkage, Black-Litterman, risk parity · 15 min

BLOCK 1 · 15 MIN

Why Diversification Works

The one free lunch in finance

Don't Put All Your Eggs in One Basket

A 2 500-year-old intuition made rigorous in 1952

The Talmud (c. 1200 BCE – 500 CE) advises:

"Let every man divide his money into three parts, and invest a third in land, a third in merchandise, and keep a third in reserve."

Harry Markowitz (1952) turned this folk wisdom into a formal mathematical framework — the birth of modern portfolio theory. He won the 1990 Nobel Memorial Prize in Economics for this work.

🎯 The core question Markowitz answered

Given a set of risky assets with known expected returns and risks, which COMBINATION of them gives:

- the **highest expected return** for a chosen risk level, or equivalently
- the **lowest risk** for a chosen expected return?

His insight: the answer depends on how the assets *co-move* (their covariance), not just on each asset alone.

A Simple Two-Asset Example

How combining assets reduces risk

Setup

- Stock A — expected return 10%, volatility 20%
- Stock B — expected return 10%, volatility 20%
- 50/50 portfolio: expected return 10% (same as each stock alone). What about risk?

General two-asset portfolio variance:

$$\sigma_p = \sqrt{w_A^2\sigma_A^2 + w_B^2\sigma_B^2 + 2\,w_A w_B\,\rho_{AB}\,\sigma_A\sigma_B}$$

Plug in σ_A = σ_B = 20%, w_A = w_B = 0.5, ρ = 0:

$$\rho = 0 \;\Rightarrow\; \sigma_p = \sqrt{0.5^2\cdot 0.20^2 + 0.5^2\cdot 0.20^2} = 14.1\%$$

⟹ 30% less risk with the SAME expected return. This is diversification's "free lunch".

How Correlation Reshapes the Picture

The three extreme cases

ρ = +1 (perfectly correlated)

$\sigma_p = 20\%$

No diversification benefit — the two stocks move identically

ρ = +0.5 (moderate correlation)

$\sigma_p \approx 17.3\%$

Modest benefit — the real-world regime for same-sector stocks

ρ = 0 (independent)

$\sigma_p = 14.1\%$

Strong benefit — classic textbook example

ρ = −0.5 (negative correlation)

$\sigma_p \approx 10.0\%$

Rare in equities; common between stocks and long-dated bonds

ρ = −1 (perfect hedge)

$\sigma_p = 0\%$

Risk fully eliminated — only a theoretical limit

BLOCK 2 · 15 MIN

Return and Risk of a Portfolio

The math — gentler than it looks

Weights and the Portfolio Setup

Notation we will use in the next 10 slides

Universe of N risky assets

- $\mu = (\mu_1, \ldots, \mu_N)^T$: expected returns vector (annualized)
- Σ $(N \times N)$: covariance matrix of returns (annualized)
- $w = (w_1, \ldots, w_N)^T$: portfolio weights, with $\sum w_i = 1$ (budget constraint)
- $\mathbf{1}$ is the N-vector of ones
- If short-selling is prohibited, we add $w_i \geq 0$ for every i

The quantities we care about

- Portfolio expected return $\mu_p = w^T\mu$ (linear in the weights — the easy part)
- Portfolio variance $\sigma^2_p = w^T\Sigma w$ (quadratic — this is why there is a theory)
- Portfolio Sharpe ratio $SR = (\mu_p - r_f) / \sigma_p$ (risk-adjusted performance)

Portfolio Expected Return

A weighted average of asset returns

Formal expression

$$\mu_p = \sum_{i=1}^{N} w_i \, \mu_i = \mathbf{w}^\top \boldsymbol{\mu}$$

Numerical example

- 3 assets with expected returns μ = (12%, 8%, 4%)
- Weights w = (0.5, 0.3, 0.2)
- μ_p = 0.5 · 12% + 0.3 · 8% + 0.2 · 4% = 8.8%

Python one-liner: `port_return = w @ mu`

The linear shape has two consequences:

- Returns can be shifted by rebalancing — no "convexity" trade-off here.
- The challenge of portfolio optimization is NOT in the return but in the risk.

The Covariance Matrix Σ

The single most important object in MPT

Elements of Σ

$$\Sigma_{ij} = \sigma_i \sigma_j \rho_{ij}, \qquad \Sigma_{ii} = \sigma_i^2$$

Properties of Σ

- **Symmetric:** $\Sigma = \Sigma^T$ (we only need to estimate $N(N+1)/2$ entries)
- **Positive semi-definite:** $w^T \Sigma w \geq 0$ for every w — variance cannot be negative
- For N = 100 assets, Σ has 5 050 unknowns. *This is why Σ is the main source of estimation error in MPT.*
- Estimated from historical returns: `Sigma = returns_df.cov() * 252`
- Rule of thumb: you need **at least 2–3x more observations than assets** for Σ to be invertible AND well-conditioned

Portfolio Variance — the Quadratic Form

This is where the geometry of MPT comes from

Formula

$$\sigma_p^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i\, w_j\, \Sigma_{ij} = \mathbf{w}^\top \Sigma\, \mathbf{w}$$

How to read the double sum

- **Diagonal terms (i = j):** $w_i^2 \cdot \sigma_i^2$ — each asset's own risk weighted by its weight squared
- **Off-diagonal terms (i ≠ j):** $w_i\, w_j \cdot \sigma_{ij}$ — pairwise co-movement
- When $\sigma_{ij} < 0$ (negative covariance), the cross-term REDUCES portfolio variance — diversification
- The quadratic structure means **volatility does NOT scale linearly with weights** — contrast with return

Python: `port_vol = np.sqrt(w @ Sigma @ w)`

BLOCK 3 · 20 MIN

The Efficient Frontier

Markowitz mean-variance framework

What is the Efficient Frontier?

The set of portfolios no rational investor would reject

Definition (Markowitz 1952)

The **efficient frontier** is the set of portfolios that offer the **maximum expected return** for every attainable level of risk (and equivalently the **minimum risk** for every expected return).

✓ On the frontier

- Not dominated by any other portfolio
- For each σ, has the highest achievable μ
- For each μ, has the lowest achievable σ
- Candidate for any risk-averse investor

✗ Below the frontier

- DOMINATED — another portfolio exists with:
  - the same σ but higher μ, OR
  - the same μ but lower σ
- No rational investor would pick these

The Optimization Problem

A convex quadratic program solved for each target return

Primal formulation

$$\min_{\mathbf{w}} \quad \mathbf{w}^\top \Sigma \mathbf{w} \quad \text{s. t.} \quad \mathbf{w}^\top \boldsymbol{\mu} = \mu^\star, \quad \mathbf{1}^\top \mathbf{w} = 1$$

Lagrangian (short-selling allowed version)

$$\mathcal{L}(\mathbf{w}, \lambda_1, \lambda_2) = \mathbf{w}^\top \Sigma \mathbf{w} - \lambda_1 (\mathbf{w}^\top \boldsymbol{\mu} - \mu^\star) - \lambda_2 (\mathbf{1}^\top \mathbf{w} - 1)$$

How to solve it in practice

- **Short-selling allowed:** closed-form via $\partial L / \partial w = 0$, a 3×3 linear system in λ_1, λ_2 and w
- **Long-only (w ≥ 0):** numerical QP — `scipy.optimize.minimize(method="SLSQP")` or `cvxpy`
- Repeat for a grid of $\mu^\star$ values to trace the frontier — about 50 points is usually enough

The Minimum-Variance Portfolio (MVP)

The leftmost — least-risky — point on the frontier

MVP weights (short-selling allowed)

$$\mathbf{w}_{\text{MVP}} = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}^{\top}\Sigma^{-1}\mathbf{1}}$$

Why use the MVP?

- **Does not require estimating μ** — the most noisy input in MPT is avoided entirely
- Very stable in-sample AND out-of-sample weights — robust to small changes in data
- Widely used as a benchmark; many real-money products are explicit minimum-variance strategies
- Clenow (2013) showed US min-vol strategies had very similar long-run returns to the market but much less drawdown

Python three-liner:

```
inv_Sigma = np.linalg.inv(Sigma)
ones = np.ones(N)
w_mvp = inv_Sigma @ ones / (ones @ inv_Sigma @ ones)
```

Geometry of the Frontier

What you will see in the (σ, μ) plane

Key properties of the (σ, μ) curve:

1. With short-selling allowed, the attainable set is a hyperbola — the efficient half is the UPPER branch (from MVP upward)

2. The MVP is at the **leftmost point** of the hyperbola — minimum σ

3. **Individual assets** (vertices of the simplex) always lie INSIDE or on the frontier — never above it

4. Adding a NEW asset to the universe can **only push the frontier up-and-left** or leave it unchanged — never down

5. Imposing constraints (e.g., $w \geq 0$) can only push the frontier **down-and-right** — constraints are costly

BLOCK 4 · 15 MIN

The Sharpe Ratio & Tangent Portfolio

Adding a risk-free asset: the Capital Market Line

The Sharpe Ratio

Reward-per-unit-of-risk

Definition (Sharpe 1966)

$$\mathrm{SR}(\mathbf{w}) = \frac{\mathbf{w}^\top \boldsymbol{\mu} - r_f}{\sqrt{\mathbf{w}^\top \Sigma \mathbf{w}}}$$

Reading the Sharpe ratio

- SR < 0 : *worse than cash* — don't hold the risky asset
- 0 < SR < 0.5 : weak risk-adjusted return — most individual stocks live here
- 0.5 < SR < 1 : decent diversified portfolios — passive indices (S&P 500 ≈ 0.5–0.7 long-run)
- 1 < SR < 2 : strong active manager territory — hedge funds targeting this range
- SR > 2 : exceptional — short-track-record or over-fitted; be skeptical

⚠ *Sharpe ratio has known limitations: assumes return normality, penalizes upside volatility the same as downside. Later work (Sortino 1994) adjusts for this.*

The Tangent Portfolio

The portfolio that maximizes the Sharpe ratio

Closed-form weights

$$\mathbf{w}_T = \frac{\Sigma^{-1}(\boldsymbol{\mu} - r_f \mathbf{1})}{\mathbf{1}^\top \Sigma^{-1}(\boldsymbol{\mu} - r_f \mathbf{1})}$$

Why is this the "right" risky portfolio?

- On a (σ, μ) chart with a risk-free asset at (0, r_f), draw the **steepest straight line** from (0, r_f) that just touches the frontier.
- The point of tangency is the Tangent Portfolio. The slope of that line = its Sharpe ratio.
- Every point on the line to the LEFT of the tangency: positive cash + long tangent portfolio.
- Every point to the RIGHT: borrow at r_f and lever up the tangent portfolio.

⟹ Every rational investor, regardless of risk aversion, holds the **same** risky portfolio. Risk preference only affects the cash share. This is the **two-fund separation theorem** (Tobin 1958).

The Capital Market Line (CML)

Every portfolio a rational investor would ever hold lies on this line

Equation of the CML

$$\mu = r_f + \frac{\mu_T - r_f}{\sigma_T}\sigma$$

Two important ideas that fall out of this:

1. With the risk-free asset available, the **efficient frontier becomes a STRAIGHT LINE** — the CML dominates the hyperbola everywhere except at the tangency

2. The CAPM (Sharpe 1964, Lintner 1965): if everyone holds the same tangent portfolio, in equilibrium that portfolio must equal the **market portfolio** — the capitalization-weighted average of all risky assets.

Two-Fund Separation in One Picture

One tangent portfolio + cash slider = every investor's portfolio

Risk-aversion levels and the corresponding allocations

Risk-free share α	Position	Investor profile
$\alpha = 1.0$	100% cash	Pure risk-averse — e.g., short-term capital preservation
$\alpha = 0.7$	30% tangent + 70% cash	Conservative — retirees, endowment fixed-income sleeve
$\alpha = 0.4$	60% tangent + 40% cash	Balanced — typical 40-year-old savings portfolio
$\alpha = 0.0$	100% tangent	Aggressive — risk-neutral, capturing full market risk premium
$\alpha = -0.5$	Borrow 50% + 150% tangent	Levered — hedge funds, risk-parity vehicles

BLOCK 5 · 30 MIN · HANDS-ON

BIST + US Live Lab

Download data, estimate inputs, build the frontier

The 7-Step Pipeline

What you will code in the Colab notebook

1. Download adjusted close prices for 8 assets (4 BIST + 4 US) via yfinance
2. Compute log returns and annualize: $\mu = 252 \cdot \text{mean}$, $\Sigma = 252 \cdot \text{cov}$
3. Simulate 10 000 random portfolios → scatter plot (σ, μ) colored by Sharpe
4. Compute Minimum-Variance Portfolio (MVP) via closed form
5. Compute Tangent (Max-Sharpe) Portfolio via closed form
6. Trace long-only efficient frontier via SLSQP
7. BONUS — Ledoit-Wolf shrinkage · Risk parity · out-of-sample backtest

💻 ***Switch to the Colab notebook now — these slides stay up for reference.***

The Lab Universe

Why we mix Turkish and US assets

GARAN.IS	Garanti Bankası	TR Banks	High vol, single-macro-factor exposure
KCHOL.IS	Koç Holding	TR Conglomerate	Lower vol — internal sector diversification
EREGL.IS	Ereğli Demir Çelik	TR Steel	Commodity-linked, FX-sensitive
THYAO.IS	Turkish Airlines	TR Aviation	Oil-price exposure, tourism cycle
SPY	S&P 500 ETF	US Equity	Benchmark for US stocks
AAPL	Apple	US Tech	Single large-cap, highest Sharpe in sample
GLD	SPDR Gold Trust	Commodity	Low correlation with equities — diversifier
TLT	20Y+ US Treasury ETF	Rate / Defensive	Negative correlation with equities in risk-off

The CRUCIAL design choice: GLD and TLT. Without them, the frontier is tight and uninteresting. These two "diversifiers" are what push the frontier up-and-left.

What You Will See in the Colab

A schematic of the key figure

The hero chart has five layered elements on the (σ, μ) plane:

① **Random cloud** — 10 000 random long-only portfolios, colored by Sharpe (viridis)

② **Individual assets** — plotted as red stars at their (σ, μ) — inside the random cloud

③ **Efficient frontier** — long-only curve (black), the upper envelope of the attainable set

④ **MVP** — the leftmost point on the frontier (navy circle)

⑤ **Tangent** — maximum-Sharpe point (gold diamond)

⑥ **Capital Market Line** — dashed line from (0, r_f) through the Tangent and beyond

BLOCK 6 · 15 MIN

Practical Limits & Extensions

Why real-world practitioners use something other than vanilla Markowitz

Three Problems with Vanilla Markowitz

And why the frontier is often called "error-maximizer"

① Estimation error

μ and Σ are ESTIMATED. The mean is especially noisy. Small errors in μ get amplified by the optimizer → extreme, concentrated weights. The optimizer loves outliers in its inputs. Michaud (1989) called MPT an "error-maximizer".

② Non-stationarity

Historical data assumes the future looks like the past. Regime changes (COVID-19 March 2020, 2022 inflation shock, Turkish lira crises) break this assumption. Both means AND covariances shift over time.

③ No investor views

Pure Markowitz is 100% data-driven. A real portfolio manager has a VIEW — "I think tech will outperform banks next year". Markowitz offers no principled way to incorporate this view into weights.

Three Modern Fixes

And their mathematical core

Ledoit-Wolf shrinkage

$$\widehat{\Sigma} = \delta \mathbf{F} + (1-\delta)\mathbf{S}$$

Shrink the sample covariance S toward a structured target F (identity, constant-correlation).

δ ∈ [0, 1] is chosen to minimize expected loss (closed-form).

✓ Cuts extreme weights
✓ sklearn.covariance.LedoitWolf
✓ Default in many funds

Black-Litterman (1992)

Bayesian update of market-implied returns with the manager's views:

π ~ equilibrium returns
P, Q ~ view matrix + magnitudes
Ω ~ view confidence

Posterior μ̃ is a weighted blend of π and the views.

✓ Stable, intuitive weights
✓ Industry standard at asset managers

Risk parity

$$w_i \cdot (\Sigma \mathbf{w})_i = w_j \cdot (\Sigma \mathbf{w})_j \quad \forall\, i, j$$

Each asset contributes EQUALLY to portfolio risk. Ignores μ — pure risk budget. Used by Bridgewater All-Weather.

✓ No μ estimation
✓ Robust across regimes
✗ Often needs leverage to match equity returns

Takeaways and Homework

What to remember and what to practice

🧠 Key concepts to remember

1. Diversification reduces risk WITHOUT reducing expected return — but only if assets are less than perfectly correlated
2. Portfolio return is LINEAR in weights ($w^T\mu$); portfolio variance is QUADRATIC ($w^T\Sigma w$) → convex optimization
3. The efficient frontier is the set of (σ, μ) pairs that dominate everything below them
4. With a risk-free asset, every rational investor holds the SAME risky mix — the Tangent (max-Sharpe) portfolio
5. In practice vanilla Markowitz is fragile — shrinkage, Black-Litterman, and risk parity are the modern fixes

💻 Your Colab practice

- Replace one BIST stock with a commodity (USO, SLV). Does the frontier shift up or down?
- Set r_f = 0.30 (real TR 1-month deposit rate). Does the Tangent portfolio still have a sensible tilt?
- Use `sklearn.covariance.LedoitWolf` instead of raw Σ — do extreme Tangent weights shrink?
- (Challenge) Backtest: fit Tangent on 2019-2022, hold through 2023-2024. Realized vs in-sample Sharpe?

Thank You

Questions?

onurpolat@hacettepe.edu.tr