

Risk, Regulation, and the Future

Day 5 of 5


Prof. Jörg Osterrieder

MSc Seminar: Digital Finance

Spring 2026

MSc Seminar: Digital Finance

The Journey So Far



Today's mission: Bring it all together.

- How to *measure* crypto risk (fat tails, not normal)
- How to *build portfolios* that include crypto
- How DeFi creates *systemic risk* through composability
- How *regulators* are responding (MiCA, SEC, future directions)

The Day DeFi Nearly Died: March 12, 2020

Black Thursday

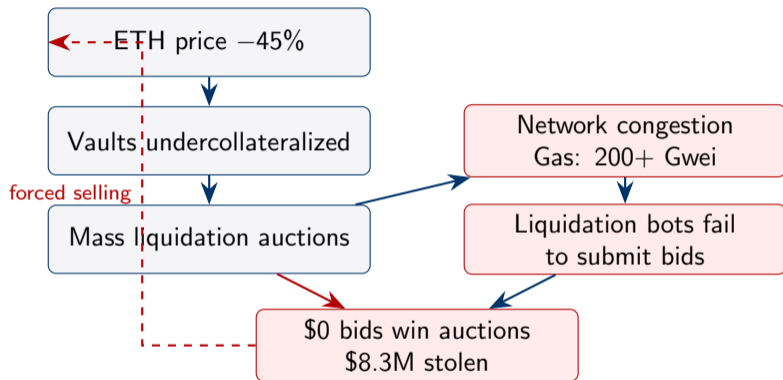
ETH price drops **45% in 24 hours** amid the COVID-19 market panic. MakerDAO's liquidation system faces its first true stress test.

- MakerDAO vaults (CDPs) require 150% collateralization
- ETH crash \Rightarrow thousands of vaults fall below threshold
- Liquidation auctions trigger simultaneously
- Ethereum network congestion: gas prices spike 10 \times
- Most liquidation bots *cannot get transactions through*

Result

A single liquidator won auctions with bids of **\$0**, seizing \$8.3M in ETH collateral for free. Vault owners lost everything.

Anatomy of a DeFi Cascade



The feedback loop—forced liquidation sales push prices lower, triggering more liquidations—is the same mechanism as traditional margin call spirals, but amplified by blockchain congestion.

DeFi Losses: The Track Record

Event	Date	Loss
MakerDAO Black Thursday	Mar 2020	\$8.3M
IRON/TITAN stablecoin collapse	Jun 2021	\$2B (market cap)
Terra/LUNA death spiral	May 2022	\$40B+
FTX collapse	Nov 2022	\$8B+ (customer funds)
Euler Finance hack	Mar 2023	\$197M

Common thread: Risk models that assumed normal distributions, adequate liquidity, and independent failures.

Today: How to model risk when tails are fat, correlations spike in crises, and DeFi protocols are interconnected.

Today: Quantify Risk, Build Portfolios, Understand Regulation

- 1 Fat-Tailed Risk Measures
- 2 Portfolio Construction with Crypto
- 3 DeFi Systemic Risk
- 4 Regulation & the Future
- 5 Hands-On: Portfolio Risk Analysis

Why Normal VaR Fails for Crypto

BTC daily returns (2015–2024):

- Kurtosis $\approx 12\text{--}15$ (normal = 3)
- Skewness ≈ -0.5 (left tail heavier)
- “5-sigma” events happen monthly, not once per millennium

Normal assumption:

- $P(|r| > 3\sigma) = 0.27\%$
- $P(|r| > 5\sigma) = 5.7 \times 10^{-5}\%$
- Predicts -10% daily return occurs once in 300 years

Crypto reality:

- $P(|r| > 3\sigma) \approx 2\%$
- Several $> 5\sigma$ events per year
- -10% days occur multiple times annually (BTC had 15 in 2022 alone)

Consequence: Normal VaR *systematically underestimates* tail risk. We need heavier-tailed distributions.

Student- t Distribution: Heavier Tails

PDF of the (standardized) Student- t distribution:

$$f(x; \nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi} \Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$$

- ν : degrees of freedom (shape parameter)
- $\nu \rightarrow \infty$: converges to the normal distribution
- $\nu = 3-5$: typical for crypto returns
- Tails decay as $|x|^{-(\nu+1)}$ (polynomial) vs. $e^{-x^2/2}$ (Gaussian)

Location-scale version for modeling returns:

$$r_t \sim \mu + s \cdot t_\nu \quad \text{where } \mu = \text{location, } s = \text{scale}$$

Fit by MLE: estimate $(\hat{\mu}, \hat{s}, \hat{\nu})$ from historical returns.

Parametric VaR with Student- t

Definition 1 (Value-at-Risk)

The α -VaR is the loss threshold exceeded with probability $1 - \alpha$:

$$\mathbb{P}(L > \text{VaR}_\alpha) = 1 - \alpha$$

Under Student- t returns:

$$\text{VaR}_\alpha = -(\mu + s \cdot t_\nu^{-1}(\alpha))$$

where $t_\nu^{-1}(\alpha)$ is the α -quantile of the Student- t distribution with ν degrees of freedom.

Comparison (BTC, $\alpha = 0.01$, daily):

Model	VaR (99%)	Actual breaches
Normal ($\sigma = 3.5\%$)	8.1%	3.2% (too many!)
Student- t ($\nu = 4$)	12.3%	1.1% (close to 1%)

Expected Shortfall (CVaR)

Definition 2 (Expected Shortfall)

The expected loss *given* that we are in the tail:

$$ES_{\alpha} = \mathbb{E}[L \mid L > VaR_{\alpha}]$$

Under Student- t :

$$ES_{\alpha} = -\mu + s \cdot \frac{f_{\nu}(q_{\alpha})}{1 - \alpha} \cdot \frac{\nu + q_{\alpha}^2}{\nu - 1}$$

where $q_{\alpha} = t_{\nu}^{-1}(\alpha) < 0$ for small α , f_{ν} is the Student- t PDF, and ES is expressed as a positive loss.

ES answers a different question:

- VaR: “What’s the worst loss at the 99% level?” \Rightarrow 12%
- ES: “If we’re in the worst 1%, *how bad* does it get?” \Rightarrow 18%
- ES captures the *severity* of tail events, not just their threshold

Why ES is Better: Coherence

Theorem 3 (Artzner et al., 1999)

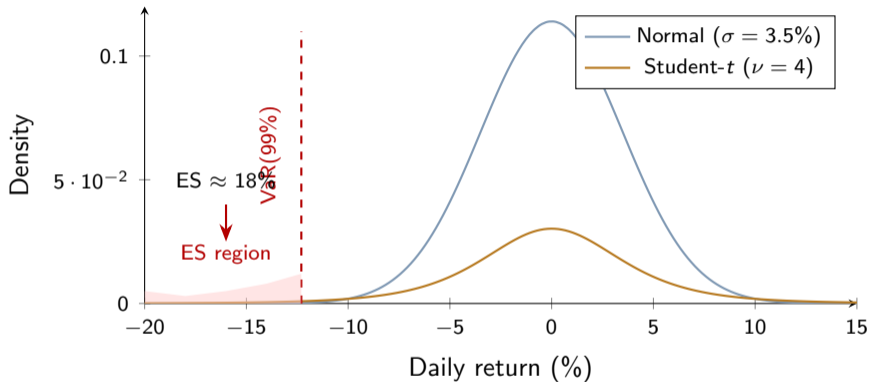
A risk measure ρ is coherent if it satisfies:

- 1 **Monotonicity:** $X \leq Y$ a.s. $\Rightarrow \rho(X) \geq \rho(Y)$
- 2 **Translation invariance:** $\rho(X + c) = \rho(X) - c$
- 3 **Positive homogeneity:** $\rho(\lambda X) = \lambda \rho(X)$
- 4 **Subadditivity:** $\rho(X + Y) \leq \rho(X) + \rho(Y)$

Subadditivity and diversification:

- VaR is *not* subadditive: combining two positions can *increase* measured VaR (contradicts diversification!)
- ES is coherent: portfolio ES \leq sum of individual ES
- Basel III adopted ES (97.5%) replacing VaR for market risk (2023)

Visualizing VaR and ES



The Student- t assigns much more probability to extreme losses. The shaded region is the ES: the average of all losses beyond VaR.

Numerical Example: BTC Risk Metrics

Data: BTC daily returns, 2020–2024 ($n = 1,461$).

Fitted Student- t : $\hat{\mu} = 0.05\%$, $\hat{\sigma} = 2.8\%$, $\hat{\nu} = 3.7$.

	VaR(95%)	VaR(99%)	ES(99%)
Normal	5.7%	8.1%	9.3%
Student- t ($\hat{\nu} = 3.7$)	6.5%	12.3%	18.1%
Historical	6.8%	13.1%	19.4%

Interpretation:

- Normal VaR(99%) = 8.1%; actual 99th percentile $\approx 13\%$
- Student- t is much closer to historical, especially in the tail
- ES reveals: when bad days happen, average loss is 18% (twice the VaR)
- A \$1M BTC position faces \$180k expected loss on bad days

Quick Question

Portfolio A has VaR of \$50k and Portfolio B has VaR of \$60k. Can their combined VaR exceed \$110k? (Hint: subadditivity)

Checkpoint

Quick Question

Portfolio A has VaR of \$50k and Portfolio B has VaR of \$60k. Can their combined VaR exceed \$110k? (Hint: subadditivity)

Answer

Yes! VaR is **not subadditive**. It is possible to construct examples where $\text{VaR}(A + B) > \text{VaR}(A) + \text{VaR}(B)$. This is one of the key reasons VaR was replaced by ES (Expected Shortfall) in Basel III—ES *is* subadditive, meaning diversification always reduces or maintains tail risk as measured by ES.

Markowitz with Crypto Assets []

Classic mean-variance optimization:

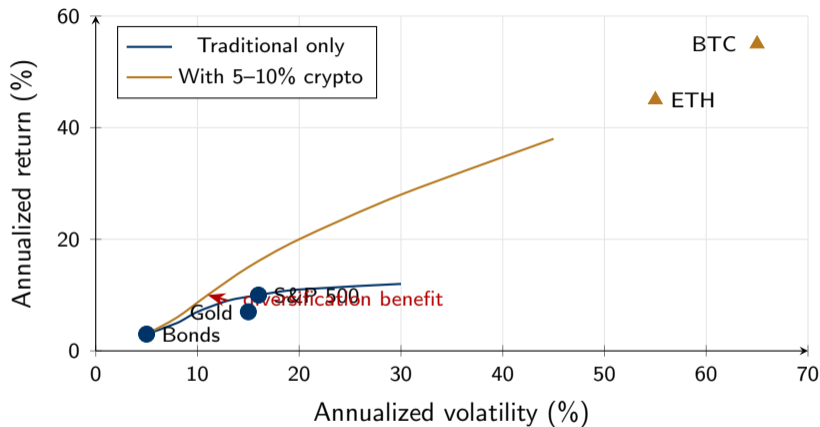
$$\min_{\mathbf{w}} \mathbf{w}'\Sigma\mathbf{w} \quad \text{s.t.} \quad \mathbf{w}'\boldsymbol{\mu} \geq \mu_{\text{target}}, \quad \mathbf{w}'\mathbf{1} = 1, \quad w_i \geq 0$$

Asset universe: S&P 500, US bonds, gold, BTC, ETH.

Asset	Ann. Return	Ann. Vol	Sharpe
S&P 500	10%	16%	0.50
US Bonds	3%	5%	0.20
Gold	7%	15%	0.33
BTC	55%	65%	0.77
ETH	45%	80%	0.50

Crypto's high Sharpe (in sample) and low correlation with traditional assets make it attractive—but estimation risk is enormous.

Efficient Frontier with Crypto Overlay



A small crypto allocation (5–10%) shifts the frontier up and left: higher return *and* better diversification—if correlations hold.

Estimation Risk and Shrinkage []

Problem: Markowitz is extremely sensitive to input estimates.

- BTC has ~ 10 years of daily data; ETH ~ 8 years
- Sample covariance matrix with 5 assets and 2,000 observations: still noisy
- Small estimation errors \Rightarrow wild portfolio swings

Ledoit-Wolf shrinkage:

$$\hat{\Sigma}_{\text{shrunk}} = (1 - \delta) \hat{\Sigma}_{\text{sample}} + \delta \hat{\Sigma}_{\text{target}}$$

- $\hat{\Sigma}_{\text{target}}$: structured estimator (e.g., diagonal or constant-correlation)
- $\delta \in [0, 1]$: optimal shrinkage intensity (data-driven)
- Reduces extreme off-diagonal elements \Rightarrow more stable portfolios

Rule: Always use shrinkage when crypto is in the asset universe. Sample covariance alone leads to absurd allocations.

CVaR Optimization: Minimize Tail Risk

Rockafellar-Uryasev (2000) formulation:

$$\begin{aligned} \min_{\mathbf{w}, \zeta} \quad & \zeta + \frac{1}{T(1-\alpha)} \sum_{t=1}^T \max(-\mathbf{w}'\mathbf{r}_t - \zeta, 0) \\ \text{s.t.} \quad & \mathbf{w}'\boldsymbol{\mu} \geq \mu_{\text{target}}, \quad \mathbf{w}'\mathbf{1} = 1, \quad w_i \geq 0 \end{aligned}$$

- ζ : auxiliary variable (approximates VaR at the optimum)
- The max term penalizes losses beyond ζ
- This is a **linear program**—solved efficiently with standard LP solvers
- Directly minimizes expected shortfall, not variance

Advantage over Markowitz: focuses on the tail. A portfolio that minimizes CVaR at 95% may accept higher variance if it avoids catastrophic losses.

Mean-Variance vs. CVaR Portfolios Under Stress

	MV Portfolio	CVaR Portfolio	Equal Weight
BTC weight	8%	5%	20%
ETH weight	4%	2%	20%
S&P 500	50%	45%	20%
Bonds	30%	42%	20%
Gold	8%	6%	20%
Annual return	12.5%	10.8%	24%
Annual vol	11.2%	9.1%	22%
Sharpe	0.93	0.96	0.95
Max drawdown	-22%	-15%	-42%
ES(99%)	4.8%	3.6%	8.5%

CVaR portfolio sacrifices return for tail protection: 30% lower max drawdown and 25% lower ES than mean-variance.

Practical Allocation Insights

- ① **5–10% crypto allocation improves Sharpe** by 0.05–0.15 for a traditional 60/40 portfolio
- ② **Correlation instability:** BTC-S&P 500 correlation was ≈ 0.1 (2017–2019) but spiked to 0.6+ during COVID and the 2022 rate-hiking cycle
- ③ **Rebalancing frequency:** monthly rebalancing works; daily is too costly in crypto
- ④ **Crypto-to-crypto diversification:** limited benefit—BTC-ETH correlation ≈ 0.85 over long horizons

The Correlation Trap

Crypto diversifies in calm markets but correlates in crises—precisely when you need diversification most. Always stress-test with crisis-era correlations.

DeFi as a Financial Network []

Traditional finance: banks, insurers, CCPs connected by interbank lending, derivatives, and payment flows.

DeFi: protocols connected by token flows, shared collateral, and composability (“money legos”).

- **Nodes:** protocols (Aave, MakerDAO, Uniswap, Lido, Compound)
- **Edges:** token flows between protocols
- **Edge weights:** total value locked (TVL) flowing through
- **Composability:** one protocol’s output is another’s input

Example Chain

User deposits ETH → Lido (receives stETH) → Aave (collateral for USDC loan) → Uniswap (provide liquidity) → earn LP tokens → use as collateral in another protocol. . .

Network Centrality Metrics

Graph $G = (V, E)$ where $V =$ protocols, $E =$ token flows.

Degree centrality:

$$C_D(v) = \frac{\text{deg}(v)}{|V| - 1}$$

How many protocols does v connect to?

Betweenness centrality:

$$C_B(v) = \sum_{s \neq v \neq t} \frac{\sigma_{st}(v)}{\sigma_{st}}$$

How often is v on the shortest path between other protocols?

Eigenvector centrality:

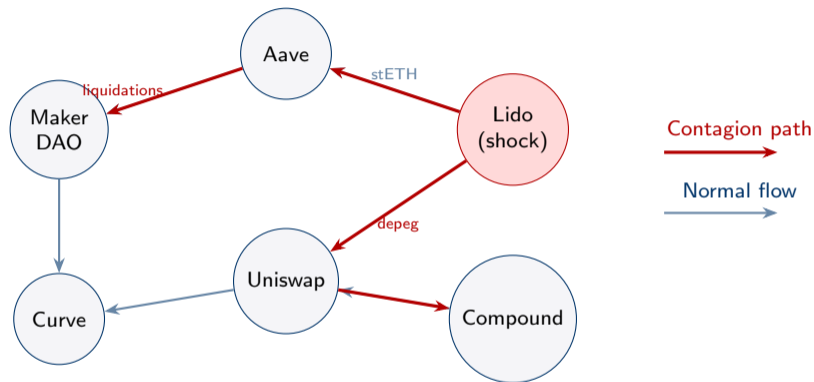
$$C_E(v) = \frac{1}{\lambda} \sum_{u \in N(v)} C_E(u)$$

Is v connected to other *important* protocols?

Key findings:

- Lido has highest eigenvector centrality (stETH is everywhere)
- USDC/USDT bridges have highest betweenness
- Failure of a high-betweenness node fragments the network

DeFi Contagion Network



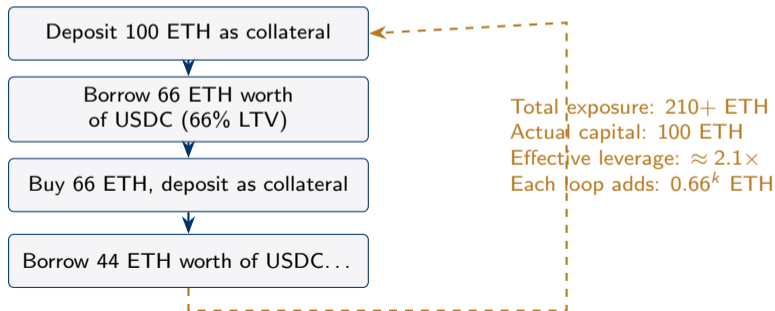
If Lido's stETH depegs from ETH, cascading liquidations propagate through Aave (collateral devaluation), Uniswap (LP losses), and downstream protocols. This is DeFi's version of "too interconnected to fail."

Recursive Borrowing and Hidden Leverage []

BIS Working Paper 1171 key findings:

- 1 Users deposit collateral, borrow stablecoins, buy more collateral, re-deposit—creating **leveraged loops**
- 2 Effective leverage can reach 3–5× through recursive positions
- 3 This leverage is *invisible* to any single protocol

Leverage loop example:



Stress Test: ETH –40% Shock

Scenario: ETH drops 40% in 24 hours (comparable to March 2020).

Protocol	Mechanism	Impact
Aave / Compound	ETH collateral devalued	\$2.5B liquidations triggered
MakerDAO	Vaults undercollateralized	DAI may depeg (upward)
Lido	stETH selling pressure	stETH/ETH ratio drops to 0.92
Uniswap	ETH/USDC pools	LPs suffer 15%+ impermanent loss
Cascading	Recursive leverage unwinds	Additional \$1B forced selling

Key insight: The \$2.5B in direct liquidations triggers an additional \$1B in forced selling from leverage unwinds, which pushes ETH lower, triggering more liquidations. The **feedback loop** amplifies the shock by 40%+.

MiCA: The EU's Comprehensive Framework []

Markets in Crypto-Assets Regulation (effective June 2023, fully applicable December 2024):

Token classification:

- 1 **E-money tokens:** pegged to single fiat (USDC, EURC) ⇒ strict reserve requirements
- 2 **Asset-referenced tokens:** pegged to baskets ⇒ reserve + governance rules
- 3 **Utility tokens & other:** ⇒ whitepaper, disclosure, conduct rules

Key requirements:

- Licensed CASPs (Crypto-Asset Service Providers)
- Reserve backing for stablecoins (1:1 for e-money tokens)
- Market abuse rules (insider trading, manipulation)
- Consumer protection (disclosure, right of withdrawal)
- Sustainability reporting for consensus mechanisms

US: Fragmented Regulation

SEC approach:

- Applies the **Howey Test** (1946): is it an “investment contract”?
- Most tokens = securities (SEC view)
- Enforcement actions against exchanges (Coinbase, Binance, Kraken)
- ETH classification: ongoing debate

CFTC approach:

- BTC and ETH are “commodities”
- Regulates crypto derivatives and futures
- Less restrictive than SEC
- Wants expanded authority

The Howey Test: An investment contract exists if there is:

- 1 An investment of money
- 2 In a common enterprise
- 3 With expectation of profits
- 4 Derived from the efforts of others

Most ICO tokens satisfy all four criteria ⇒ unregistered securities.

Tokenization: \$10B+ On-Chain, Projected \$2–4T by 2030

Real-world assets (RWA) on blockchain:

Already tokenized:

- US Treasuries: \$2.5B+ on-chain (Ondo, Franklin Templeton)
- Private credit: \$9B+ (Maple, Goldfinch, Centrifuge)
- Real estate: \$500M+ (RealT, Lofty)
- Carbon credits: \$300M+ (Toucan, KlimaDAO)

Benefits:

- 24/7 settlement (vs. T+1 or T+2)
- Fractional ownership (\$10 real estate)
- Programmable compliance (auto-transfer restrictions)
- Transparent cap tables
- Composability with DeFi

BlackRock BUIDL Fund

Tokenized US Treasury fund launched March 2024. \$500M+ AUM by end of 2024.
Institutional validation of on-chain finance.

The Road Ahead

Near-term (2025–2027):

- 1 **AI + DeFi agents:** autonomous agents that manage portfolios, arbitrage, and provide liquidity
- 2 **ZK proofs for compliance:** prove you're not on a sanctions list without revealing identity
- 3 **Account abstraction:** better UX, social recovery, gasless transactions

Longer-term:

- 1 **CBDC integration:** central bank money on DeFi rails
- 2 **Cross-chain interoperability:** seamless asset movement
- 3 **Quantum threats:** current cryptography vulnerable; migration to post-quantum schemes needed by 2030s
- 4 **Institutional DeFi:** regulated pools, KYC-gated protocols

The big picture: Traditional finance is moving on-chain. DeFi is becoming regulated. They're converging—and the students who understand both will shape the result.

Hands-On Session

Portfolio Risk Analysis and Stress Testing

Python notebooks on course platform

Step 1: Load Multi-Asset Returns

Data setup

```
import pandas as pd
import numpy as np
from scipy import stats

# Load daily returns: BTC, ETH, SPY, TLT, GLD (2020-2024)
returns = pd.read_csv('multi_asset_returns.csv',
                     parse_dates=['date'],
                     index_col='date')

print(returns.describe())
print(f"Kurtosis:\n{returns.kurtosis()}")
```

Tasks:

- 1 Compute summary statistics (mean, std, skew, kurtosis)
- 2 Plot return distributions with normal overlay

Step 2: Fit Student- t and Compute VaR/ES

Student- t fit

```
# Fit Student-t to BTC returns
nu, loc, scale = stats.t.fit(returns['BTC'])
print(f"BTC: nu={nu:.1f}, loc={loc:.4f}, scale={scale:.4f}")

# Parametric VaR and ES
alpha = 0.01
VaR_t = -(loc + scale * stats.t.ppf(alpha, nu))
# ES via integration
ES_t = -loc + scale * stats.t.pdf(stats.t.ppf(alpha, nu), nu) \
      / (1 - alpha) * (nu + stats.t.ppf(alpha, nu)**2) / (nu - 1)
```

Compare: Normal VaR vs. Student- t VaR vs. historical VaR for each asset. Which model fits best?

Step 3: Correlation Analysis

Rolling correlations

```
# 60-day rolling correlation: BTC vs SPY
rolling_corr = returns['BTC'].rolling(60).corr(returns['SPY'])
rolling_corr.plot(title='BTC-SPY Rolling Correlation (60d)')
```

Tasks:

- 1 Plot the full correlation matrix (heatmap)
- 2 Plot rolling 60-day correlations for all crypto-traditional pairs
- 3 Identify the March 2020 and November 2022 correlation spikes
- 4 Implement Ledoit-Wolf shrinkage on the covariance matrix

Key finding: Correlations that are 0.1 in calm markets jump to 0.6+ during crises.

Step 4: Efficient Frontier Construction

Mean-variance optimization

```
from scipy.optimize import minimize

def portfolio_vol(w, cov):
    return np.sqrt(w @ cov @ w)

def neg_sharpe(w, mu, cov, rf=0.04/252):
    ret = w @ mu
    vol = portfolio_vol(w, cov)
    return -(ret - rf) / vol

# Constraints: weights sum to 1, all >= 0
constraints = {'type': 'eq', 'fun': lambda w: np.sum(w) - 1}
bounds = [(0, 1) for _ in range(n_assets)]
```

Step 5: CVaR Optimization

Linear programming formulation

```
from scipy.optimize import linprog

def cvar_optimize(returns, alpha=0.05, target_return=0.0003):
    T, n = returns.shape
    # Variables: [w (n), zeta (1), u (T)]
    # Minimize: zeta + 1/(T*(1-alpha)) * sum(u)
    # Subject to: u_t >= -w'r_t - zeta, u_t >= 0
    # w'mu >= target, sum(w) = 1, w >= 0
    ... # see notebook for full implementation
```

Tasks:

- 1 Implement CVaR optimization (95% level)
- 2 Compare CVaR-optimal weights with MV-optimal weights
- 3 How much crypto does each approach allocate?

Step 6: Stress Test — March 2020 Revisited

Use actual data from Feb 19 – Mar 23, 2020:

- 1 Compute portfolio returns for MV, CVaR, and equal-weight portfolios during the crash
- 2 Track daily drawdown paths
- 3 Compute realized VaR and ES over the stress period
- 4 Which portfolio recovered fastest?

	MV	CVaR	Equal Wt
Max drawdown	-18%	-12%	-35%
Days to recovery	45	32	80
Worst daily loss	-6.2%	-4.8%	-12.1%

The CVaR portfolio's tail-risk focus pays off exactly when it matters most.

Step 7: Full Results Summary

Compile your findings:

- ① **Risk metrics table:** VaR/ES for each asset under each model
- ② **Correlation regime chart:** calm vs. crisis correlations
- ③ **Efficient frontier:** with/without crypto, with/without shrinkage
- ④ **Portfolio weights:** MV vs. CVaR allocations
- ⑤ **Stress test:** drawdown paths and recovery

Write a 1-page summary:

- What is the optimal crypto allocation for a conservative investor?
- How does the allocation change if you use crisis-era correlations?
- What are the limits of these models?

Final Task: Executive Summary

Write a 3-sentence executive summary

Summarize your portfolio analysis for a **non-technical investment committee**.

Your summary should answer:

- Which portfolio (MV, CVaR, or equal-weight) do you recommend?
- What is the key risk the committee should know about?
- What is the single most important caveat or limitation?

Why this matters

Translating quantitative results into actionable language is a core skill for any finance professional. This scaffolds LO5: creating professional-grade analysis.

Debate: Institutional Crypto Exposure Limits

Motion

“Regulated financial institutions should be prohibited from holding more than 2% of assets under management in crypto-assets.”

For the motion:

- Crypto volatility is 4× equities
- Tail risk is extreme (daily -20% events)
- Custody and operational risk
- Correlation instability undermines diversification
- Regulatory uncertainty

Against the motion:

- Small allocation (5–10%) improves Sharpe
- Institutions can manage risk with derivatives
- Blanket limits are paternalistic
- Crypto market is maturing (ETFs, custody solutions)
- Missing the upside is also a risk

Case Discussion: Tokenized Real Estate

Scenario: A Swiss pension fund considers investing 1% of its CHF 5B portfolio in tokenized commercial real estate on Ethereum. **Discussion questions:**

- 1 What are the **advantages** over traditional real estate funds? (Liquidity, fractional ownership, transparency)
- 2 What are the **risks**? (Smart contract, regulatory, oracle, blockchain downtime)
- 3 How would you **stress test** this allocation? (ETH crash + real estate downturn simultaneously)
- 4 What **regulatory approvals** are needed under Swiss law and MiCA?
- 5 Is the 24/7 settlement a *feature* or a *risk* for a pension fund? (Weekend liquidations, flash crashes)

Case Study Assignment & Peer Review

Assignment (due in 2 weeks):

- 1 Choose a DeFi protocol or crypto-asset class
- 2 Apply risk analysis techniques from this week:
 - Fit fat-tailed distributions, compute VaR/ES
 - Analyze network position (if DeFi)
 - Backtest a simple strategy with proper methodology
- 3 Write a 10-page report with code appendix

Peer review process:

- Each student reviews two classmates' reports
- Structured rubric provided (methodology, rigor, presentation)
- Reviews are graded for quality and constructiveness
- Final grade: 60% report + 20% peer review + 20% presentation

Thank You!

What you now have:

- ① **Day 1–2:** Crypto pricing, DeFi mechanics, AMMs, options
- ② **Day 3:** Game theory of consensus, fee mechanism design, MEV
- ③ **Day 4:** ML pipeline for crypto—features, models, SHAP, realistic backtesting
- ④ **Day 5:** Fat-tailed risk, portfolio construction, systemic risk, regulation

The toolkit: You can now quantify crypto risk, build data-driven trading models, understand DeFi incentives, and navigate the regulatory landscape. These skills are in high demand across finance, fintech, and central banking.

Questions? • Office hours: by appointment • joerg.osterrieder@usi.ch

References I

- [1] European Commission. *Regulation (EU) 2023/1114 on Markets in Crypto-Assets (MiCA)*. Official Journal of the European Union. 2023.
- [2] Lioba Heimbach and Luying Huang. *DeFi Leverage*. BIS Working Paper 1171. Bank for International Settlements, 2024.
- [3] Olivier Ledoit and Michael Wolf. “A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices”. In: *Journal of Multivariate Analysis* 88.2 (2004), pp. 365–411.