Day 5: Risk, Regulation, and the Future of Digital Finance

Network Fragility, Tail Risk, CBDCs, and What Comes Next

Prof. Dr. Jörg Osterrieder

PhD Seminar Series: Digital Finance Research

2026

PhD Seminar Series: Digital Finance Research

The Week So Far

Day 1: Crypto Derivatives

- Jump-diffusion models (Merton, Kou)
- Stochastic volatility (Heston, Bates)
- Equilibrium pricing & Fourier methods

Day 2: DeFi Mathematics

- CFMM geometry [2]
- Impermanent loss, concentrated liquidity
- Lending protocol dynamics

Day 3: Blockchain Economics

- Game theory & mechanism design
- MEV and sandwich attacks [13]
- Transaction fee mechanisms

Day 4: ML & Microstructure

- Kyle model, VPIN [14]
- Reinforcement learning [12]
- Transformer models for LOB forecasting

The $2.57 Billion Wash Trading Problem

Chainalysis 2024 Findings

- **$2.57B** in wash trading volume detected across DEXs
- 67% concentrated in low-liquidity token pairs
- Motivated by token airdrops and artificial volume inflation
- Detection: same-entity round-trip transactions within blocks

Detection Heuristic

Flag address a if:

$$\sum_{t \in \mathcal{T}_a} \mathbb{1}[\text{counterparty}(t) \in \mathcal{C}(a)] > \tau$$

where $\mathcal{C}(a)$ is the cluster of addresses linked to a via on-chain heuristics.

Pseudonymity $\neq$ anonymity, but enforcement is hard at scale.

FBI “Operation Token Mirrors”

The Operation (2024)

- FBI created **NexFundAI** — a fake ERC-20 token
- Targeted market makers engaged in wash trading
- **18 individuals and entities charged**
- **$25M** in fraudulent market manipulation exposed
- First-ever use of undercover token by law enforcement

Scheme Mechanics

1. Market makers offered “volume services”
2. Coordinated buy/sell orders to inflate price
3. Created illusion of organic demand
4. Retail investors bought at inflated prices
5. Insiders dumped — classic pump-and-dump

Lesson: Traditional fraud, new medium.

The Challenge: Regulating a Permissionless System

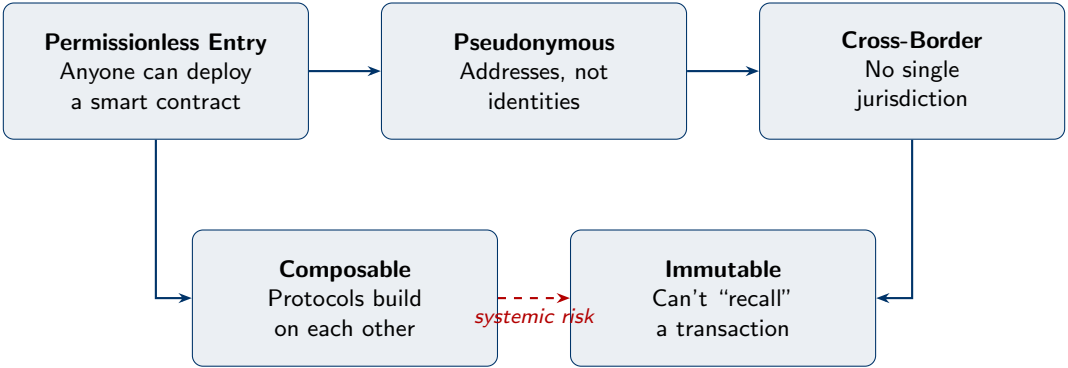

"You cannot subpoena a smart contract." — regulatory reality

What We Need: Three Pillars

Risk Measurement	Systemic Analysis	Regulatory Frameworks
• Fat-tailed distributions	• Network topology	• CBDC design
• Extreme Value Theory	• Contagion cascades	• Stablecoin regulation
• Copula dependence	• Clearing mechanisms	• RegTech platforms
• VaR/CVaR under non-normality	• Leverage amplification	• Cross-border coordination

Goal: Bring mathematical rigor to digital finance governance.

Today's Roadmap

1. Network Theory and Systemic Risk
2. Portfolio Theory for Digital Assets
3. CBDCs and Stablecoins
4. Applications and the Future

DeFi as a Directed Graph

Definition 1

The **DeFi network** is a directed weighted graph $G = (V, E, w)$ where:

- $V = \{v_1, \ldots, v_n\}$: protocols (Aave, Compound, MakerDAO, ...)
- $E \subseteq V \times V$: capital flows between protocols
- $w : E \to \mathbb{R}^+$: weight = TVL (total value locked) flowing on edge

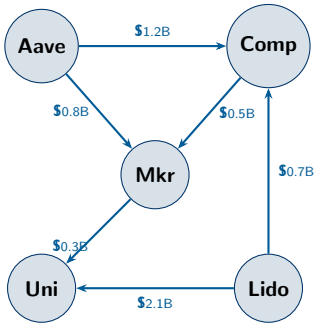

Composability creates hidden interconnections [7].

Adjacency Matrix and Weighted Degree Centrality

Definition 2

The **weighted adjacency matrix** $\boldsymbol{A} \in \mathbb{R}_{\geq 0}^{n \times n}$:

$$A_{ij} = w(v_i, v_j) = \text{TVL flowing from protocol } i \text{ to protocol } j$$

Centrality Measures

Weighted in-degree (how much capital flows *into* protocol i):

$$k_i^{\text{in}} = \sum_{j=1}^{n} A_{ji}$$

Weighted out-degree (exposure *from* protocol i):

$$k_i^{\text{out}} = \sum_{j=1}^{n} A_{ij}$$

Betweenness Centrality: The Bottleneck Nodes

Definition 3

The **betweenness centrality** of node v measures how often v lies on shortest paths between other nodes:

$$c_B(v) = \sum_{\substack{s,t \in V \\ s \neq v \neq t}} \frac{\sigma_{st}(v)}{\sigma_{st}} \tag{1}$$

where σ_{st} is the number of shortest paths from s to t, and $\sigma_{st}(v)$ is the number of those paths passing through v.

DeFi Interpretation

- High $c_B(v) \Rightarrow$ protocol v is a **critical intermediary**
- Wrapped tokens (e.g., wETH, stETH) often have highest betweenness
- Failure of high-betweenness node **fragments** the network
- Normalized: $\hat{c}_B(v) = \frac{c_B(v)}{(n-1)(n-2)/2}$

Eigenvector Centrality: Importance from Neighbors

Definition 4

The **eigenvector centrality** $\boldsymbol{x} \in \mathbb{R}^n$ satisfies:

$$\boldsymbol{A}\boldsymbol{x} = \lambda_1 \boldsymbol{x}, \qquad x_i = \frac{1}{\lambda_1} \sum_{j=1}^{n} A_{ij}\, x_j \tag{2}$$

where λ_1 is the largest eigenvalue of $\boldsymbol{A}$ (Perron–Frobenius).

Properties

- A node is important if its *neighbors* are important — recursive
- Perron–Frobenius theorem: for $A_{ij} \geq 0$ and $\boldsymbol{A}$ irreducible, $\exists$ unique $\boldsymbol{x} > 0$ with $\|\boldsymbol{x}\|_1 = 1$
- Convergent power iteration: $\boldsymbol{x}^{(k+1)} = \frac{\boldsymbol{A}\boldsymbol{x}^{(k)}}{\|\boldsymbol{A}\boldsymbol{x}^{(k)}\|_1}$
- Google's PageRank is a variant: $\boldsymbol{x} = \alpha \boldsymbol{A}\boldsymbol{x} + (1-\alpha)\frac{\mathbf{1}}{n}$

Application: Identifying systemically important DeFi protocols ("DeFi-SIFIs").

Eisenberg–Noe (2001): Clearing Vector Framework

Setting [ ]

- n financial institutions (or DeFi protocols) with interlinked obligations
- $\mathbf{\Pi} \in \mathbb{R}_{\geq 0}^{n \times n}$: relative liability matrix, Π_{ij} = fraction of i's total obligations owed to j
- $\bar{\boldsymbol{p}} \in \mathbb{R}_{+}^{n}$: vector of total nominal obligations
- $\boldsymbol{e} \in \mathbb{R}^{n}$: exogenous operating cash flow (external assets)

Key Insight

Each institution's ability to pay depends on what it *receives* from others, which depends on their ability to pay — a **fixed-point problem**.

Central question: Given a network of obligations and an external shock, *which institutions default and how much does each pay?*

The Clearing Vector p^*

Definition 5

A **clearing vector** $\boldsymbol{p}^* \in \mathbb{R}_+^n$ satisfies, for each node i:

$$p_i^* = \min\left(e_i + \sum_{j=1}^{n} \Pi_{ji}\, p_j^*, \;\; \bar{p}_i \right) \tag{3}$$

Interpretation:

- Total resources of i: external assets e_i plus payments received $\sum_j \Pi_{ji} p_j^*$
- If resources $\geq$ obligations: pay in full ($p_i^* = \bar{p}_i$)
- If resources $<$ obligations: pay what you can (proportional default)

In Vector Form

$$\boldsymbol{p}^* = \min\big(\boldsymbol{e} + \boldsymbol{\Pi}^\top \boldsymbol{p}^*,\; \bar{\boldsymbol{p}}\big) \qquad \Longleftrightarrow \qquad \boldsymbol{p}^* = \Phi(\boldsymbol{p}^*), \quad \Phi(\boldsymbol{p}) := \min\big(\boldsymbol{e} + \boldsymbol{\Pi}^\top \boldsymbol{p},\; \bar{\boldsymbol{p}}\big)$$

Fixed-Point Existence and Computation

Theorem 6 (Eisenberg–Noe, 2001)

The map $\Phi : [0, \bar{\boldsymbol{p}}] \to [0, \bar{\boldsymbol{p}}]$ *is:*

1. **Monotone**: $\boldsymbol{p} \leq \boldsymbol{q} \implies \Phi(\boldsymbol{p}) \leq \Phi(\boldsymbol{q})$
2. **Self-mapping**: $\Phi([0, \bar{\boldsymbol{p}}]) \subseteq [0, \bar{\boldsymbol{p}}]$

By Tarski's fixed-point theorem, $\exists$ *a* **greatest** *clearing vector* $\boldsymbol{p}^*$ *and a* **least** *clearing vector* $\boldsymbol{p}_*$.

Computation: Monotone Iteration ("Fictitious Default Algorithm")

1. Initialize $\boldsymbol{p}^{(0)} = \bar{\boldsymbol{p}}$ (everyone pays in full)
2. Iterate: $\boldsymbol{p}^{(k+1)} = \Phi(\boldsymbol{p}^{(k)})$
3. Convergence in at most n steps (one new defaulter per step)

Complexity: $O(n^2)$ per iteration, $O(n^3)$ worst case. Practical for $n \leq 10^4$.

Adapting Eisenberg–Noe to DeFi

Key Differences from Traditional Finance

- **Collateral ratios** replace nominal obligations
- Obligations are **over-collateralized**: $\mathrm{CR}_i = \frac{\mathrm{Collateral}_i}{\mathrm{Debt}_i} > 1$
- Liquidation is **automatic**: triggered when $\mathrm{CR}_i < \mathrm{CR}_{\min}$
- No bankruptcy courts — smart contracts enforce pro-rata distribution

Modified Clearing Condition

For DeFi protocol i with collateral C_i and debt D_i:

$$p_i^* = \min\left(C_i(S) + \sum_j \Pi_{ji}\, p_j^*,\ \ D_i \right)$$

where $C_i(S)$ depends on asset price S — **price-mediated contagion**.

Collateral is typically ETH or stablecoins whose value can drop, coupling the clearing problem to market dynamics.

Threshold Cascades in DeFi

Definition 7

Protocol i **fails** (enters liquidation) when:

$$\mathrm{CR}_i = \frac{C_i(S)}{D_i - R_i(\boldsymbol{p})} < \theta_i \tag{4}$$

where $R_i(\boldsymbol{p}) = \sum_j \Pi_{ji} p_j$ is total recovery from other protocols and θ_i is the liquidation threshold (e.g., $\theta_i = 1.5$ for Aave).

Cascade Dynamics

1. **Shock**: ETH price drops $S \to S' = S(1 - \delta)$
2. **Round 1**: Protocols with $\mathrm{CR}_i(S') < \theta_i$ are liquidated
3. **Fire sale**: Liquidations push price further: $S' \to S'' < S'$
4. **Round 2**: More protocols breach threshold at S''
5. **Repeat** until a new fixed point or total collapse

Cascade Simulation: ETH Price Shock

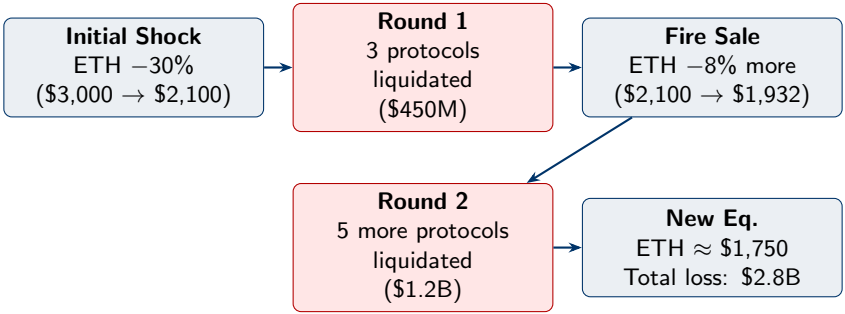

Amplification Factor

$$\alpha = \frac{\text{Total loss}}{\text{Initial shock value}} = \frac{\$2.8\text{B}}{\$0.9\text{B} \times 0.30} \approx 10.4\times$$

Leverage creates a **loss multiplier** — small shocks cascade into large losses.

BIS WP 1062: DeFi Is More Fragile than TradFi

Key Findings [ ]

- DeFi network is **scale-free**: degree distribution follows a power law $P(k) \sim k^{-\gamma}$ with $\gamma \approx 2.3$
- **Few hubs** (Aave, Compound, MakerDAO) concentrate $> 60\%$ of TVL flows
- Average path length $\ell \approx 2.1$ — shocks propagate in ≤ 3 hops
- Clustering coefficient $\mathcal{C} \approx 0.41$ — high local interconnection

Fragility vs. Traditional Finance

- TradFi: regulated capital buffers, lender of last resort, circuit breakers
- DeFi: **no** capital requirements, **no** central bank backstop, **no** trading halts
- Composability is a feature *and* a systemic vulnerability
- Recursive leverage amplifies losses by 3–10$\times$ [8]

Fat Tails: Crypto Returns Are NOT Gaussian

Empirical Facts

- BTC daily returns: kurtosis $\hat{\kappa} \approx 12$–16 (Gaussian: $\kappa = 3$)
- ETH: $\hat{\kappa} \approx 18$–25
- Skewness: often negative during crashes
- **Tail events** are 5–50$\times$ more frequent than Gaussian predicts

Consequence

Standard risk measures **fail**:

$$\mathbb{P}[|r_t| > 4\sigma] \approx \begin{cases} 6.3 \times 10^{-5} & \text{Gaussian} \\ 0.01\text{–}0.05 & \text{empirical} \end{cases}$$

A "once in 43 years" event under Gaussian assumptions occurs **several times per year** in crypto markets [11].

Quantile–Quantile Plot: BTC vs. Normal

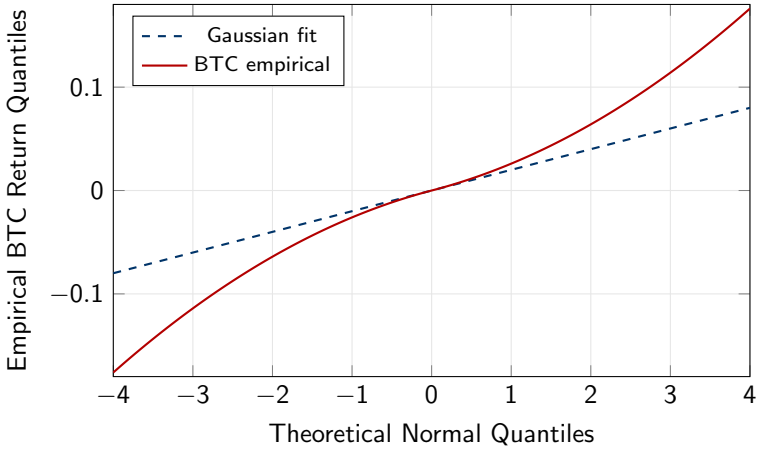

Concave left tail, convex right tail ⇒ **heavier tails than Gaussian**. Departure grows at extremes: precisely where risk management matters most.

Student-t Distribution: A First Improvement

Definition 8

The standardized Student-t distribution with ν degrees of freedom:

$$f(x;\nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-(\nu+1)/2} \tag{5}$$

Properties

- **Tail decay**: polynomial $\sim |x|^{-(\nu+1)}$ vs. Gaussian $\sim e^{-x^2/2}$
- Kurtosis $= 3 + \frac{6}{\nu-4}$ for $\nu > 4$ — controlled by single parameter
- BTC: $\hat{\nu} \approx 3$–5 (very heavy tails, infinite 4th moment if $\nu \leq 4$)
- Variance exists only if $\nu > 2$; mean only if $\nu > 1$

Better than Gaussian, but symmetric — cannot capture skewness. Use skewed-t (Hansen, 1994) or Generalized Hyperbolic for asymmetry.

Extreme Value Theory: Block Maxima Approach

Motivation

We care about **rare, extreme losses** — the tail of the tail. EVT provides a principled framework without assuming a full distributional model.

Definition 9 (Block Maxima)

Divide return series into m blocks of size n. Let $M_j = \max(r_1^{(j)}, \ldots, r_n^{(j)})$. By the Fisher–Tippett–Gnedenko theorem:

$$\frac{M_j - b_n}{a_n} \xrightarrow{d} H_\xi \qquad \text{as } n \to \infty$$

where H_ξ is one of three types determined by the **shape parameter** ξ.

Three Domains of Attraction

- $\xi > 0$: **Fréchet** — heavy tails (power law). *Crypto lives here.*
- $\xi = 0$: **Gumbel** — light tails (exponential decay). *Gaussian, etc.*

The Generalized Extreme Value (GEV) Distribution

Definition 10

The GEV distribution unifies all three types:

$$F(x;\mu,\sigma,\xi) = \exp\left\{-\left[1+\xi\left(\frac{x-\mu}{\sigma}\right)\right]^{-1/\xi}\right\} \tag{6}$$

defined on $\{x : 1+\xi(x-\mu)/\sigma > 0\}$, with:

- $\mu \in \mathbb{R}$: location parameter
- $\sigma > 0$: scale parameter
- $\xi \in \mathbb{R}$: shape (tail index)

Estimation for Crypto

- BTC weekly block maxima: $\hat{\xi} \approx 0.2$–0.4 (heavy tails confirmed)
- MLE: $\hat{\boldsymbol{\theta}} = \arg\max_{\mu,\sigma,\xi} \sum_{j=1}^{m} \log f_{\text{GEV}}(M_j;\mu,\sigma,\xi)$
- Profile likelihood CIs preferred over asymptotic normality for ξ

Value-at-Risk from the GEV Distribution

Definition 11

The **Value-at-Risk** at confidence level α is the α-quantile of the loss distribution:

$$\mathbb{P}[L > \text{VaR}_\alpha] = 1 - \alpha$$

Theorem 12 (GEV Quantile)

Inverting the GEV CDF for $\xi \neq 0$:

$$VaR_\alpha = \mu + \frac{\sigma}{\xi}\left[(-\log\alpha)^{-\xi} - 1\right] \tag{7}$$

Numerical Example (BTC, weekly)

With $\hat{\mu} = 0.02$, $\hat{\sigma} = 0.04$, $\hat{\xi} = 0.30$:

$$\text{VaR}_{0.99} = 0.02 + \frac{0.04}{0.30}\left[(-\log 0.99)^{-0.30} - 1\right] \approx 0.185 \quad \text{(18.5\% weekly loss)}$$

Expected Shortfall (CVaR)

Definition 13

Expected Shortfall (Conditional VaR) at level α:

$$\mathrm{ES}_\alpha = \mathbb{E}\big[L \mid L > \mathrm{VaR}_\alpha\big] = \frac{1}{1-\alpha}\int_\alpha^1 \mathrm{VaR}_u \,\mathrm{d}u$$

Theorem 14 (ES from GEV, $\xi < 1$)

$$ES_\alpha = VaR_\alpha + \frac{\sigma + \xi\,(VaR_\alpha - \mu)}{1-\xi} \tag{8}$$

Why ES over VaR?

- VaR answers: "How bad can it get (at level α)?"
- ES answers: "*Given* we exceeded VaR, *how bad on average*?"
- ES is **coherent** (subadditive): $\mathrm{ES}(X+Y) \le \mathrm{ES}(X) + \mathrm{ES}(Y)$
- VaR is **not** subadditive — diversification can "increase" VaR

Copulas: Beyond Correlation

Theorem 15 (Sklar, 1959)

For any joint CDF $F(x_1, \ldots, x_d)$ *with marginals* $F_1, \ldots, F_d$, $\exists$ *a copula* $C : [0,1]^d \to [0,1]$ *such that:*

$$F(x_1, \ldots, x_d) = C\big(F_1(x_1), \ldots, F_d(x_d)\big) \tag{9}$$

If $F_1, \ldots, F_d$ *are continuous,* C *is unique.*

Why Copulas for Crypto?

- Correlation ρ captures *linear* dependence only
- Copulas model the full dependence structure, including **tail dependence**
- Crypto assets can be nearly uncorrelated in normal markets but **crash together** in stress — tail dependence $\neq$ correlation
- **Lower tail dependence**: $\lambda_L = \lim_{u \to 0^+} \mathbb{P}[U_2 \leq u \mid U_1 \leq u]$
- **Upper tail dependence**: $\lambda_U = \lim_{u \to 1^-} \mathbb{P}[U_2 > u \mid U_1 > u]$

Clayton Copula: Lower Tail Dependence

Definition 16

The bivariate **Clayton copula** with parameter $\theta > 0$:

$$C_\theta(u, v) = \left(u^{-\theta} + v^{-\theta} - 1\right)^{-1/\theta} \tag{10}$$

Properties

- **Lower tail dependence**: $\lambda_L = 2^{-1/\theta} > 0$
- **Upper tail dependence**: $\lambda_U = 0$ (asymmetric)
- As $\theta \to 0^+$: independence ($C = uv$)
- As $\theta \to \infty$: comonotonicity ($C = \min(u, v)$)

Crypto Application

BTC-ETH pair: $\hat{\theta} \approx 1.8 \Rightarrow \lambda_L = 2^{-1/1.8} \approx 0.68$.

Interpretation: Given BTC is in its worst 1% of outcomes, there is a 68% probability ETH is also in its worst 1%.

Student-t Copula: Symmetric Tail Dependence

Definition 17

The bivariate t-**copula** with ν d.f. and correlation ρ:

$$C_{\nu,\rho}(u,v) = t_{\nu,\rho}\big(t_\nu^{-1}(u),\, t_\nu^{-1}(v)\big) \tag{11}$$

where t_ν^{-1} is the univariate t-quantile function and $t_{\nu,\rho}$ is the bivariate t CDF with correlation ρ.

Tail Dependence

$$\lambda_L = \lambda_U = 2\, t_{\nu+1}\left(-\sqrt{\frac{(\nu+1)(1-\rho)}{1+\rho}}\right) > 0 \quad \text{for all } \nu < \infty$$

Crypto Implication

- Cryptos **crash together AND rally together** — symmetric tail dependence
- BTC-ETH: $\hat{\nu} \approx 4$, $\hat{\rho} \approx 0.75 \Rightarrow \lambda \approx 0.40$
- Gaussian copula: $\lambda_L = \lambda_U = 0$ always — **dangerously wrong**

Portfolio Optimization: Minimize CVaR, Not Variance

Classical Mean-Variance (Markowitz)

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

Problem: Assumes Gaussian returns. Covariance $\boldsymbol{\Sigma}$ misses tail risk.

CVaR Optimization (Rockafellar & Uryasev, 2000)

$$\min_{\boldsymbol{w}, z} \quad z + \frac{1}{T(1-\alpha)} \sum_{t=1}^{T} \max\left(-\boldsymbol{w}^\top \boldsymbol{r}_t - z,\ 0\right) \tag{12}$$

subject to $\boldsymbol{w}^\top \boldsymbol{\mu} \geq r_{\text{target}}$, $\boldsymbol{w}^\top \mathbf{1} = 1$, $\boldsymbol{w} \geq 0$.

- This is a **linear program** — efficient even for large portfolios
- Directly uses *historical scenarios* — no distributional assumption
- For crypto: CVaR-optimal portfolios hold 30–50% less in high-tail-risk assets

DeFi Leverage Amplification

BIS WP 1171: Recursive Borrowing [ ]

DeFi enables **recursive leverage**:

1. Deposit 1 ETH as collateral on Aave (LTV ratio $\ell = 0.75$)
2. Borrow ℓ ETH worth of stablecoins
3. Buy ETH with stablecoins, re-deposit as collateral
4. Repeat k times

Effective Leverage after k Rounds

$$L_k = \sum_{i=0}^{k} \ell^i = \frac{1-\ell^{k+1}}{1-\ell} \xrightarrow{k\to\infty} \frac{1}{1-\ell} \tag{13}$$

With $\ell = 0.75$: $L_\infty = 4\times$. With $\ell = 0.825$ (some protocols): $L_\infty \approx 5.7\times$.

Systemic Risk

- Effective leverage is **invisible** to any single protocol

CBDC Landscape: 91% of Central Banks Exploring

Global Status (BIS 2024) [ ]

- **134** countries actively exploring CBDCs
- **91%** of central banks in some stage of CBDC work
- **3** fully launched: Bahamas, Jamaica, Nigeria
- **36** in pilot phase (incl. China, India, Brazil)
- **EU**: digital euro legislation advanced

Motivations

- Financial inclusion
- Payment efficiency
- Monetary policy transmission
- Counter private stablecoins
- Programmable money

Key tension: efficiency vs. financial stability (bank disintermediation risk).

SF Fed DSGE Model of CBDC

Model Structure [ ]

Three-agent New Keynesian DSGE:

- **Households**: choose consumption C_t, deposits D_t, CBDC M_t, bonds B_t
- **Banks**: accept deposits, make loans, face reserve requirements
- **Central bank**: sets policy rate i_t, CBDC interest rate i_t^m

Household Portfolio Choice

Households maximize:

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Big[u(C_t) + \chi \cdot v(D_t, M_t) - h(N_t) \Big]$$

where $v(D_t, M_t)$ captures liquidity services from deposits and CBDC:

$$v(D_t, M_t) = \left[\alpha D_t^{\frac{\varepsilon-1}{\varepsilon}} + (1-\alpha) M_t^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}$$

CES aggregator with elasticity ε between deposits and CBDC.

Optimal CBDC Interest Rate

Theorem 18 (Jiang & Zhu, 2024)

In the calibrated steady state, the welfare-maximizing CBDC rate satisfies:

$$i_t^{m*} \approx \max\big(0\%,\ i_t - \Delta^*\big) \qquad \textit{with } \Delta^* \approx 1\% \tag{14}$$

Intuition

- CBDC rate **too high**: households shift deposits → CBDC, banks lose funding, credit contracts
- CBDC rate **too low**: CBDC adoption is minimal, benefits of digital currency unrealized
- **Goldilocks:** CBDC rate ≈ policy rate -1%
- The 1% spread compensates banks for their intermediation role

First-order condition balances marginal liquidity benefit of CBDC against marginal cost of bank disintermediation.

Bank Disintermediation: Quantitative Results

Calibrated Predictions [ ]

Scenario	Deposit Outflow	Lending Reduction
CBDC rate $= 0\%$	-1.8%	-0.4%
CBDC rate $= i_t - 1\%$	-3.2%	-0.9%
CBDC rate $= i_t$	-8.7%	-3.1%
CBDC rate $= i_t + 0.5\%$	-15.4%	-6.8%

Bank Profit Channel

Bank j profit: $\pi_j = (r_j^L - r_j^D)D_j - \text{costs}$. As D_j falls:

$$\frac{\partial \pi_j}{\partial M_t} = -(r_j^L - r_j^D)\frac{\partial D_j}{\partial M_t} < 0$$

Banks respond by: (1) raising deposit rates, (2) reducing lending, (3) seeking wholesale funding.

Retail vs. Wholesale CBDC

Retail CBDC

- Direct access for households and firms
- Replaces cash (partially)
- Requires identity infrastructure (KYC)
- **Risk:** bank disintermediation
- **Benefit:** financial inclusion, direct monetary policy transmission
- Privacy concerns: central bank sees all transactions

Wholesale CBDC

- Restricted to financial institutions
- Upgrades interbank settlement
- Minimal impact on deposits
- **Risk:** limited direct benefit to public
- **Benefit:** faster cross-border settlement, DvP for securities
- Easier regulatory compliance

Hybrid models (two-tier): central bank issues, commercial banks distribute.
Most advanced pilots (China's e-CNY, digital euro) follow the hybrid approach.

Stablecoin Mechanism Design

Collateralized (USDC, USDT)

- 1:1 reserve backing (fiat, treasuries)
- Redemption: burn token $\rightarrow$ receive \$1
- Peg stability: arbitrage

$$P > 1 \Rightarrow \text{mint \& sell}$$

$$P < 1 \Rightarrow \text{buy \& redeem}$$

- Risk: reserve quality, counterparty

Algorithmic (historical)

- No external collateral
- Dual-token: stable + governance
- Mint/burn mechanism:

$$P < 1 \Rightarrow \text{burn stable, mint gov}$$

- Relies on **confidence**: governance token must retain value
- Risk: reflexive death spiral

Crypto-Collateralized (DAI)

Over-collateralized: $\text{CR} = \frac{\text{ETH collateral}}{\text{DAI minted}} \geq 1.5$. Liquidation auction if CR falls below threshold.

Terra/Luna: The Death Spiral

Mechanism Failure (May 2022)

1. UST (stable) pegged via mint/burn with LUNA (governance)
2. To redeem \$1 UST: burn UST, receive \$1 worth of LUNA
3. **Trigger:** Large UST sell-off $\Rightarrow P_{\text{UST}} < 1$
4. Arbitrageurs burn UST $\rightarrow$ mint LUNA $\rightarrow$ sell LUNA
5. LUNA supply $\uparrow\uparrow \Rightarrow$ LUNA price $\downarrow\downarrow$
6. Need *more* LUNA to redeem \$1 $\Rightarrow$ faster dilution

The Reflexivity Trap

$$\underbrace{P_{\text{UST}} \downarrow}_{\text{depeg}} \Rightarrow \underbrace{\text{LUNA minted} \uparrow}_{\text{redemptions}} \Rightarrow \underbrace{P_{\text{LUNA}} \downarrow}_{\text{dilution}} \Rightarrow \underbrace{\text{Confidence} \downarrow}_{\text{panic}} \Rightarrow P_{\text{UST}} \downarrow\downarrow$$

Result: \$40B market cap $\rightarrow$ \$0 in 5 days. No external collateral = no floor.

Regulatory Landscape: MiCA, US, and Beyond

EU: MiCA (2024–25)

- Markets in Crypto-Assets Regulation
- **Stablecoins**: 1:1 reserve + liquid assets, regular audits
- **Significant stablecoins**: stricter capital, EBA supervision
- **CASPs**: licensed, AML/KYC, governance requirements
- First comprehensive crypto framework globally

United States

- Fragmented: SEC, CFTC, state regulators
- Stablecoin bills: reserve requirements, issuer licensing
- SEC enforcement actions (Howey test for tokens)
- Bipartisan momentum on stablecoin legislation (2025)
- SAB 121 repeal: banks can custody crypto

Key Regulatory Principles Emerging

- "Same activity, same risk, same regulation" [4]
- Technology-neutral: regulate function, not form

Applications & Industry

Tokenization · Institutional Adoption · RegTech

Real-World Asset Tokenization

The RWA Thesis [ ]

Represent off-chain assets (bonds, real estate, commodities) as on-chain tokens.

Market Size (2025)

- ¿**$10B** tokenized government securities
- ¿**$3B** tokenized private credit
- ¿**$1B** tokenized commodities (gold)
- Growth: ∼300% YoY in 2024

Benefits

- 24/7 settlement (T+0 vs. T+2)
- Fractional ownership
- Programmable compliance
- Transparent audit trail
- Global liquidity access

Open questions: Legal enforceability, oracle reliability, cross-chain interoperability, and the "last mile" problem (connecting on-chain tokens to off-chain legal rights).

Institutional Adoption: BUIDL and Beyond

BlackRock BUIDL Fund (2024)

- **BlackRock USD Institutional Digital Liquidity Fund** — on Ethereum
- AUM: $> \$500$M within first year
- Invests in US Treasuries, repos, cash
- Token = share of money market fund (ERC-20)
- Daily yield accrual, on-chain transfers, 24/7 settlement

Franklin Templeton On-Chain Money Market

- **FOBXX**: $> \$400$M AUM, on Stellar and Polygon
- Each token represents one share of the fund
- Blockchain as *secondary* record-keeping (transfer agent remains primary)
- Uses smart contracts for dividend distribution

Signal: Largest asset managers tokenizing — not "if" but "how fast."

Institutional Portfolio: 60/30/10 Framework

Multi-Layer Portfolio Construction

Core (60%): BTC + ETH + tokenized Treasuries — low turnover, strategic

Satellite (30%): DeFi yield (Aave, Lido), staking, LP positions — active

Buffer (10%): Stablecoins + CBDC — liquidity, margin, rebalancing

Risk Constraints

$$\mathrm{ES}_{0.975}(\boldsymbol{w}) \leq \bar{L}, \qquad \sum_{i \in \mathrm{DeFi}} w_i \leq 0.30, \qquad \text{max leverage} \leq 2\times$$

CVaR constraint ensures tail risk budget is respected even under fat tails.

RegTech Landscape for Digital Assets

State of the Field [ ]

Systematic survey: **41 commercial** + **28 academic** RegTech platforms.

Commercial Categories

- **On-chain analytics:** Chainalysis, Elliptic, TRM Labs
- **Transaction monitoring:** AML/CFT compliance
- **Tax reporting:** automated gain/loss
- **Smart contract audit:** formal verification
- **Identity:** decentralized KYC (zk-based)

Academic Contributions

- Graph-based address clustering
- ML for anomaly detection
- Privacy-preserving compliance
- Cross-chain tracing
- NLP for regulatory text analysis

Gap identified: limited integration between on-chain analytics and traditional financial surveillance systems.

Global RegTech Market Growth

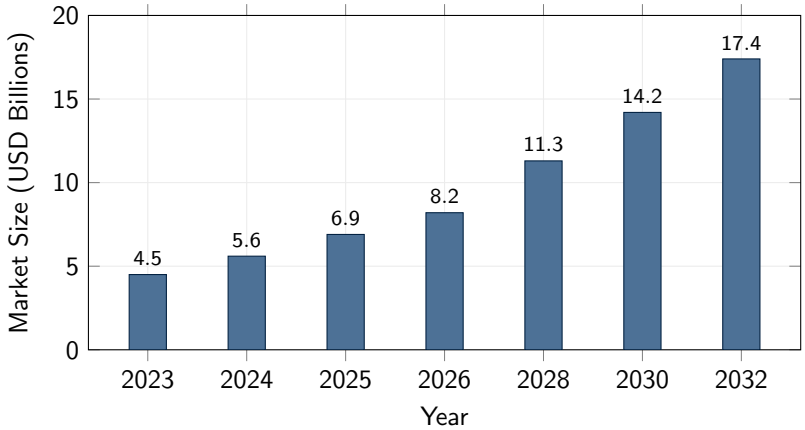

CAGR $\approx$ 16.2% (2023–2032). Driven by MiCA compliance deadlines, increasing institutional participation, and cross-border regulatory coordination.

Discussion: Cross-Chain MEV and AI Agents in DeFi

Cross-Chain MEV [ ]

- MEV extraction across L1s, L2s, bridges
- Atomic cross-chain arbitrage
- Shared sequencers: new MEV attack surface
- **Open problem:** optimal cross-chain MEV redistribution mechanism

AI Agents as DeFi Participants

- Autonomous portfolio rebalancing
- RL-based liquidation bots [12]
- AI $\times$ MEV: smarter extraction
- **Regulatory Q:** Who is liable when an AI agent causes a flash crash?

Research Frontiers

- Can mechanism design make MEV extraction welfare-neutral?
- How do AI agents change market microstructure equilibria?
- Composable game theory framework [1]

Discussion: Privacy (ZK Proofs) vs. Compliance

The Fundamental Tension

ZK-Based Selective Disclosure

Prove statement $\phi(x)$ is true **without revealing** x:

$$\text{Prover} \xrightarrow{\pi} \text{Verifier}: \quad \mathbb{P}[\text{Verify}(\phi, \pi) = 1 \mid \phi(x) = \text{true}] = 1$$

- "My balance $>$ \$10,000" without revealing exact balance
- "I am not on OFAC sanctions list" without revealing identity
- "Transaction source is compliant" without revealing the source

Open problem: ZK compliance that satisfies both FATF Travel Rule *and* user privacy expectations. Active

Seminar Wrap-Up: Key Takeaways

What We Covered

- **D1** Blockchain foundations, consensus
- **D2** AMMs, microstructure, MEV
- **D3** Derivatives, stochastic models
- **D4** ML, deep learning, RL
- **D5** Risk, regulation, CBDCs

Unifying Themes

- Math $\rightarrow$ mechanism $\rightarrow$ market
- Tail risk demands new tools
- Composability: feature + vulnerability
- Regulation must be code-aware

Research Proposal Guidance

- Pick a **specific** research question
- Ground it in **real data** (on-chain, exchange)
- Use **rigorous methodology** (EVT, copulas, network models, ML)
- Connect to **economic theory**
- Proposal due: [see syllabus]

Strong Topics

- Cascade modeling with real DeFi data
- Copula-based crypto portfolio risk
- CBDC impact on bank lending (empirical)
- MEV redistribution mechanism design

Thank You

Questions, Discussion, Research Ideas

Prof. Dr. Jörg Osterrieder
joerg.osterrieder@...

References I

[1] Guillermo Angeris et al. *A Composable Game-Theoretic Framework for Blockchains*. 2025. arXiv: 2504.18214 [cs.GT].

[2] Guillermo Angeris et al. "The Geometry of Constant Function Market Makers". In: *Proceedings of the 25th ACM Conference on Economics and Computation (EC 2024)*. 2024.

[3] Raphael Auer et al. *Real-World Asset Tokenization*. 2025. arXiv: 2503.01111 [q-fin.GN].

[4] Bank for International Settlements. *Cryptocurrencies and Decentralised Finance*. BIS Paper 156. Bank for International Settlements, 2024.

[5] Sebastian Doerr et al. *SoK: Web3 RegTech*. 2025. arXiv: 2512.24888 [cs.CR].

[6] Larry Eisenberg and Thomas H. Noe. "Systemic Risk in Financial Systems". In: *Management Science* 47.2 (2001), pp. 236–249.

[7] Lewis Gudgeon et al. *Systemic Fragility in Decentralized Markets*. BIS Working Paper 1062. Bank for International Settlements, 2023.

[8] Lioba Heimbach and Luying Huang. *DeFi Leverage*. BIS Working Paper 1171. Bank for International Settlements, 2024.

[9] International Monetary Fund. *CBDC: Progress and Further Considerations*. Tech. rep. International Monetary Fund, 2024.

[10] Janet Jiang and Shaowen Zhu. *A Macroeconomic Model of CBDC*. Working Paper 2024-11. Federal Reserve Bank of San Francisco, 2024.

[11] Shimon Kogan, Igor Makarov, Marina Niessner, and Antoinette Schoar. "Are Cryptos Different? Evidence from Retail Trading". In: *Journal of Financial Economics* 159 (2024), p. 103893.

[12] Petter N. Kolm and Gordon Ritter. *Reinforcement Learning in Financial Decision Making*. 2025. arXiv: 2512.10913 [q-fin.CP].

[13] Baran Oz, Simon Sui, and Jean Tirole. "Maximal Extractable Value and Allocative Inefficiency". In: *Journal of Financial Economics* (2025).

[14] Justin Sirignano and Rama Cont. "Deep Limit Order Book Forecasting". In: *Quantitative Finance* (2025).