Day 2: DeFi

The Mathematics of Automated Market Makers and Lending Protocols

Prof. Jörg Osterrieder

PhD Seminar – Digital Finance Research

2026

Recap Day 1 + Today's Roadmap

Yesterday – Crypto Derivatives:

- Jump-diffusions: Merton, Kou
- Stochastic vol: Heston, Bates
- Fourier pricing via characteristic functions
- Equilibrium with liquidity risk
- Kou best-fits BTC; Bates best-fits ETH

Today – DeFi Mathematics:

1. **CFMM Framework:** axiomatic AMM theory
2. **Uniswap V2 & V3:** constant product, concentrated liquidity
3. **Impermanent Loss:** rigorous derivation, V3 amplification
4. **DeFi Lending:** interest rate models, optimal control
5. **Applications:** empirical evidence, risk management

Key shift

Day 1: pricing derivatives on crypto assets (option theory).
Day 2: crypto assets *as* financial infrastructure (mechanism design + stochastic control).

What is DeFi?

Decentralized Finance (DeFi): financial services executed by smart contracts on public blockchains, without centralized intermediaries.

Scale (as of 2025):

- Total Value Locked (TVL): \$80–120B
- Daily DEX volume: \$5–15B
- Lending TVL: \$30–40B
- Stablecoin supply: $>$ \$150B

Key protocol categories:

- **DEXs:** Uniswap, Curve, Balancer
- **Lending:** Aave, Compound, Morpho
- **Derivatives:** dYdX, Synthetix, GMX
- **Yield:** Lido (staking), Yearn

How is this different?

- **Permissionless:** anyone can provide liquidity or borrow
- **Transparent:** all state on-chain, verifiable
- **Composable:** protocols can be composed ("money LEGOs")
- **Non-custodial:** users retain asset control

References:

- [7] – textbook overview
- [4] – academic treatment

Uniswap in 5 Minutes

Uniswap = the dominant decentralized exchange. No order book. No market makers (in the traditional sense). Instead: an **automated market maker** (AMM).

Mechanism:

1. A **liquidity pool** holds reserves of two tokens: (x, y) – e.g., (ETH, USDC)
2. **Liquidity providers (LPs)** deposit equal value of both tokens
3. **Traders** swap one token for the other; the pool adjusts reserves
4. **Pricing rule:** the pool maintains a **constant product invariant**

Constant product rule (Uniswap V2)

$$x \cdot y = k \qquad \text{(invariant before fees)}$$

A trader selling Δx of token X receives Δy of token Y such that:

$$(x + \Delta x)(y - \Delta y) = k \quad \Longrightarrow \quad \Delta y = \frac{y \cdot \Delta x}{x + \Delta x}$$

The Constant Product Formula: $x \cdot y = k$

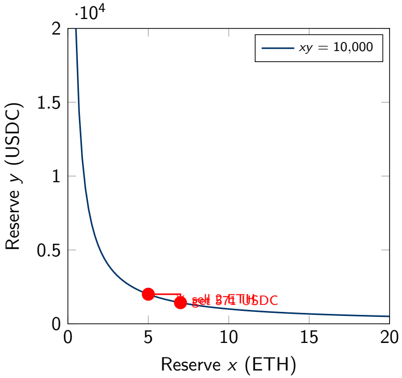

Properties:

- Hyperbolic curve in (x, y) space
- **Price:** slope of tangent line

$$P = \frac{y}{x} \quad \text{(marginal price)}$$

- Reserves can never reach zero (asymptotic)
- **Slippage:** larger trades move further along the curve → worse effective price
- Price impact $\propto \Delta x / x$

With fee γ (e.g., $\gamma = 0.997$):

$$(x + \gamma \Delta x)(y - \Delta y) = k$$

Fee $= (1 - \gamma)\Delta x$ accrues to LPs.

Providing Liquidity: How LPs Enter

Uniswap V2: LP deposits equal *value* of both tokens at the current pool price.

Example: pool has $(x_0, y_0) = (100 \text{ ETH}, 200{,}000 \text{ USDC})$, so $P_0 = 2{,}000$ USDC/ETH.

- LP deposits 10 ETH + 20,000 USDC (10% of pool)
- Receives LP tokens representing 10% share of pool
- New reserves: $(110, 220{,}000)$, new $k = 24{,}200{,}000$

LP's position:

- Pro-rata claim on pool reserves at withdrawal time
- Earns trading fees proportional to share
- **But:** reserves change as price moves – the LP's token mix shifts

The catch: Impermanent Loss

If price moves (in either direction), the LP's portfolio underperforms a simple buy-and-hold ("HODL") strategy. This is **impermanent loss** – the central risk of LP positions.

Impermanent Loss: LP Underperforms HODL

Setup: LP deposits at price P_0. Price moves to P_T. Define $r = P_T/P_0$.

HODL portfolio value:

$$V_{\text{HODL}} = x_0 P_T + y_0 = x_0 P_0 \left(\frac{1+r}{2}\right) \cdot 2$$

(proportional to $1 + r$)

LP portfolio value: At price P_T, pool rebalances to (x_T, y_T) with $x_T y_T = k$ and $y_T/x_T = P_T$:

$$V_{\text{LP}} = 2\sqrt{k \cdot P_T} \propto \sqrt{r}$$

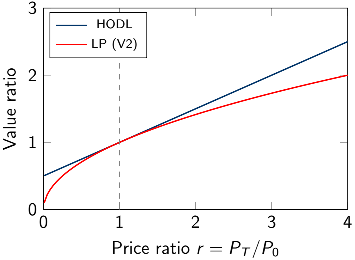

Key insight: $\sqrt{r} \leq (1 + r)/2$ for all $r > 0$ (AM-GM inequality). Equality only at $r = 1$. The LP position is a **concave function** of price – always underperforms the linear HODL.

Today: Formalize This Mathematically

Questions we will answer rigorously:

1. What is the **axiomatic foundation** of CFMMs? (Angeris et al.)
2. How does **Uniswap V3** change the game with concentrated liquidity?
3. What is the **exact formula** for impermanent loss in V2 and V3?
4. Why is an LP position equivalent to a **short straddle**?
5. How do **DeFi lending rates** emerge from stochastic control?
6. What does the **empirical evidence** say? (Lehar-Parlour, BIS)

Mathematical tools:

- Convex optimization (CFMM framework)
- Itô calculus (IL as gamma exposure)
- Stochastic control / HJB equations (lending rate optimization)
- Empirical finance (19M transaction dataset)

Outline

1. CFMM Framework
2. Impermanent Loss
3. DeFi Lending
4. Applications & Industry

AMM Primitives: Reserves and Trading Functions

Setup [1]: an AMM for n tokens.

Definitions

- **Reserves:** $R = (R_1, \dots, R_n) \in \mathbb{R}_+^n$ – current token holdings in the pool
- **Trade:** $\Delta = (\Delta_1, \dots, \Delta_n) \in \mathbb{R}^n$ – net token flow ($\Delta_i > 0$: tokens into pool)
- **Trading function:** $\varphi : \mathbb{R}_+^n \to \mathbb{R}$ – defines the AMM's pricing rule
- **Fee:** $\gamma \in (0, 1]$ – fraction of input tokens that count toward reserves

Trade feasibility: a trade Δ is valid if and only if

$$\varphi(R + \gamma\Delta^+ - \Delta^-) \geq \varphi(R)$$

where $\Delta^+ = \max(\Delta, 0)$ and $\Delta^- = \max(-\Delta, 0)$.

Interpretation: after the trade, the pool's "invariant" must not decrease. The fee $(1 - \gamma)$ on inputs means less is counted toward the invariant.

Trade Feasibility and the No-Arbitrage Condition

Simplified (no fees, $\gamma = 1$):

$$\varphi(R + \Delta) \geq \varphi(R) \qquad \text{for all feasible trades } \Delta$$

Key properties of well-defined CFMMs:

1. **Conservation:** the AMM always “wins” – φ is non-decreasing along trades
2. **No free lunch:** $\Delta \leq 0$ (all tokens leave the pool) $\implies \varphi(R + \Delta) < \varphi(R)$
3. **Prices:** at equilibrium, the trade saturates the constraint:

$$\varphi(R + \Delta) = \varphi(R)$$

4. **Marginal price** of token i in terms of token j:

$$P_{ij} = \frac{\partial \varphi / \partial R_i}{\partial \varphi / \partial R_j}$$

The surface $\{ R \in \mathbb{R}^n_+ : \varphi(R) = c \}$ is the **trading set** – all reachable reserve states from initial R_0 with $\varphi(R_0) = c$.

Angeris et al. (EC 2024): Axiomatic CFMM Theory

Main contribution [1]: a unified, rigorous framework for all CFMMs.

Axioms for a valid CFMM trading function φ:

1. $\varphi : \mathbb{R}^n_+ \to \mathbb{R}$ is **concave**
2. φ is **nondecreasing**: $R' \geq R \implies \varphi(R') \geq \varphi(R)$
3. φ is **1-homogeneous**: $\varphi(\lambda R) = \lambda\,\varphi(R)$ for all $\lambda > 0$

Why these axioms?

- **Concavity** → convex feasible trade set → unique optimal trade
- **Nondecreasing** → adding tokens never hurts (no free extraction)
- **1-homogeneous** → doubling liquidity doubles the trading capacity, prices unchanged

Canonical form theorem

Every CFMM satisfying Axioms 1–3 has a unique (up to scaling) **canonical trading function**: concave, 1-homogeneous, nondecreasing.

Theorem: Uniqueness of the Canonical Trading Function

Theorem (Angeris et al. 2024)

Let $\mathcal{T} \subseteq \mathbb{R}^n$ be a CFMM trading set satisfying standard regularity conditions (closed, convex, comprehensive). Then there exists a unique (up to positive scaling) function $\varphi : \mathbb{R}^n_+ \to \mathbb{R}$ that is:

1. Concave
2. 1-homogeneous: $\varphi(\lambda R) = \lambda \varphi(R)$
3. Nondecreasing

such that the feasible trades are exactly $\mathcal{T} = \{\Delta : \varphi(R + \Delta) \geq \varphi(R)\}$.

Proof idea: Construct φ as the **support function** of the convex set of "net transfer vectors" that the AMM accepts. 1-homogeneity follows from the conical structure of the trading set. Concavity from convexity of the complement.

Significance: any CFMM, no matter how exotic, can be described by a single canonical function. This unifies Uniswap, Curve, Balancer, and novel designs under one theory.

Special Cases: Constant Product, Sum, Mean

CFMM	$\varphi(R)$	**Price** P_{12}	**Protocol**
Constant product	$\sqrt{R_1 R_2}$	R_2/R_1	Uniswap V2
Constant sum	$R_1 + R_2$	1 (fixed)	Mento
Constant mean	$R_1^{w_1} R_2^{w_2}$	$\frac{w_1}{w_2}\frac{R_2}{R_1}$	Balancer
StableSwap	$A\sum R_i + \prod R_i$	≈ 1 near peg	Curve

Constant product: $\varphi = \sqrt{xy}$ is concave, 1-homogeneous, nondecreasing. Satisfies all axioms.

Constant sum: $\varphi = x + y$ gives fixed exchange rate (no price discovery). Reserves deplete to zero for one token if price deviates.

Constant mean (Balancer): weighted generalization. Weights (w_1, w_2) with $w_1 + w_2 = 1$ determine the LP's exposure ratio. $w_1 = w_2 = 0.5$ recovers Uniswap.

StableSwap (Curve): hybrid between constant-sum and constant-product. Amplification parameter A controls curvature near the peg.

Price in a CFMM: Marginal Rate of Substitution

Marginal price of token i in terms of token j:

$$\boxed{P_{ij}(R) = \frac{\partial\varphi/\partial R_i}{\partial\varphi/\partial R_j}}$$

Derivation: Consider an infinitesimal trade $(\mathrm{d}R_i, \mathrm{d}R_j)$ along the level set $\varphi(R) = c$:

$$\frac{\partial\varphi}{\partial R_i}\,\mathrm{d}R_i + \frac{\partial\varphi}{\partial R_j}\,\mathrm{d}R_j = 0 \implies \frac{\mathrm{d}R_j}{\mathrm{d}R_i} = -\frac{\partial\varphi/\partial R_i}{\partial\varphi/\partial R_j}$$

For Uniswap V2: $\varphi = \sqrt{xy}$, $\frac{\partial\varphi}{\partial x} = \frac{1}{2}\sqrt{y/x}$, $\frac{\partial\varphi}{\partial y} = \frac{1}{2}\sqrt{x/y}$, so:

$$P_{xy} = \frac{\sqrt{y/x}}{\sqrt{x/y}} = \frac{y}{x}$$

Economic interpretation: the price is the slope of the indifference curve (level set of φ). Exactly analogous to the MRS in microeconomic theory – the AMM's "preferences" over reserve compositions are encoded in φ.

Multi-Asset CFMM as Convex Optimization

Optimal routing problem: a trader with utility $U(\Delta)$ over net token transfers $\Delta \in \mathbb{R}^n$.

Trader's optimization (Angeris et al. 2021)

$$\max_{\Delta \in \mathbb{R}^n} \; U(\Delta) \quad \text{subject to} \quad \varphi(R + \gamma\Delta^+ - \Delta^-) \geq \varphi(R)$$

Key properties:

- φ concave $\implies$ feasible set is convex $\implies$ **convex optimization** (if U is concave)
- Unique optimal trade (under strict concavity)
- **KKT conditions** yield equilibrium prices as Lagrange multipliers

Multi-pool routing: with m pools (different CFMMs), the joint problem is:

$$\max_{\Delta^1,\dots,\Delta^m} \; U\Big(\sum_{j=1}^{m} \Delta^j\Big) \quad \text{s.t.} \quad \varphi_j(R_j + \gamma_j(\Delta^j)^+ - (\Delta^j)^-) \geq \varphi_j(R_j), \; \forall j$$

Still convex. This is the foundation of DEX aggregators (1inch, CoW Swap, Uniswap X).

Optimal Trade with Fees

Explicit form for two-token constant-product pool with fee $1-\gamma$:
Trader sells $\Delta_x > 0$ tokens of X to buy Δ_y tokens of Y:

$$\max_{\Delta_x > 0} U(\Delta_x, \Delta_y) \quad \text{s.t.} \quad (x + \gamma\Delta_x)(y - \Delta_y) = xy$$

Solution for the constraint:

$$\Delta_y = \frac{\gamma\, y\, \Delta_x}{x + \gamma\, \Delta_x}$$

Effective execution price:

$$P_{\text{eff}} = \frac{\Delta_y}{\Delta_x} = \frac{\gamma\, y}{x + \gamma\, \Delta_x}$$

Price impact:

$$\text{Slippage} = 1 - \frac{P_{\text{eff}}}{P_{\text{mid}}} = 1 - \frac{\gamma\, x}{x + \gamma\, \Delta_x} \approx \frac{\Delta_x}{x} \quad \text{for small trades}$$

Key insight

Price impact is $O(\Delta_x / x)$ – proportional to trade size relative to reserves. This is the AMM analogue of Kyle's lambda [10].

Uniswap V3: Concentrated Liquidity

Innovation (May 2021): LPs choose a **price range** $[p_a, p_b]$ to concentrate their liquidity.

V2: liquidity spread across $(0, \infty)$

- Most liquidity unused (price far from extremes)
- Capital inefficient

V3: liquidity concentrated in $[p_a, p_b]$

- Only active when $P \in [p_a, p_b]$
- Much higher capital efficiency
- Behaves like a V2 pool with "virtual reserves"

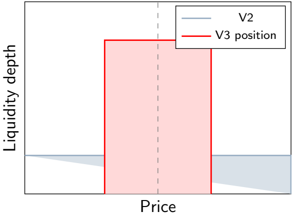

Capital efficiency: a concentrated position in $[p_a, p_b]$ provides the same depth as a V2 position with $\frac{\sqrt{p_b}-\sqrt{p_a}}{\sqrt{p_b/p_a}-1}$ times more capital.

V3: Virtual Reserves and the Liquidity Parameter

Within a tick range $[p_a, p_b]$: V3 behaves like V2 but with shifted ("virtual") reserves.

V3 invariant within $[p_a, p_b]$

$$\left(x + \frac{L}{\sqrt{p_b}}\right)\left(y + L\sqrt{p_a}\right) = L^2$$

Liquidity: $L = \sqrt{k}$ measures the pool's depth at the active price.

Reserves as function of price $P \in [p_a, p_b]$:

$$x(P) = L\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{p_b}}\right)$$

$$y(P) = L\left(\sqrt{P} - \sqrt{p_a}\right)$$

At the boundaries:

- $P = p_a$: position is 100% token X (no Y), $x = L(1/\sqrt{p_a} - 1/\sqrt{p_b})$
- $P = p_b$: position is 100% token Y (no X), $y = L(\sqrt{p_b} - \sqrt{p_a})$

Capital Efficiency: Concentrated Position as Leverage

Concentration factor: for a V3 position in $[p_a, p_b]$ around current price P:

$$\text{Efficiency ratio} = \frac{L_{V3}}{L_{V2}} = \frac{\sqrt{P}}{\sqrt{P} - \sqrt{p_a}} \cdot \frac{\sqrt{p_b}}{\sqrt{p_b} - \sqrt{P}} \cdot \frac{1}{\sqrt{p_b / p_a}}$$

Example: ETH/USDC at $P = 2000$, range $[1800, 2200]$:

- $p_a/P = 0.9$, $p_b/P = 1.1$ ($\pm 10\%$ range)
- Efficiency $\approx 10\times$ – a \$10K V3 position provides same depth as \$100K V2

Leverage analogy:

- Concentrated LP $\approx$ leveraged LP exposure
- Narrower range → higher leverage → higher fee income **and** higher IL
- In the limit $p_a \to p_b$: infinite leverage, LP collapses to a limit order

Key trade-off

Concentration amplifies **both** fee revenue and impermanent loss. The LP must choose the range optimally.

V2 vs. V3: Comparison

Feature	**Uniswap V2**	**Uniswap V3**
Liquidity range	$(0, \infty)$	$[p_a, p_b]$ (LP choice)
Capital efficiency	$1\times$	$\sim 4000\times$ max
Fungibility	Fungible LP tokens (ERC-20)	Non-fungible (NFT)
Fee tiers	0.3% only	0.01%, 0.05%, 0.3%, 1%
LP management	Passive (set and forget)	Active (rebalance range)
IL severity	Standard	Amplified (per unit capital)
Pricing formula	$xy = k$	$L^2 = (x + a)(y + b)$ per tick
Launched	May 2020	May 2021

Liquidity distribution (V3):

- Liquidity is **piecewise constant** across ticks
- Each tick: 1 basis point (0.01%) price increment
- Aggregate liquidity at price P: $L(P) = \sum_{j:\, P \in [p_a^j, p_b^j]} L_j$
- Creates a "liquidity landscape" – concentrated near current price

Outline

1. CFMM Framework
2. Impermanent Loss
3. DeFi Lending
4. Applications & Industry

Impermanent Loss: The LP Payoff is Concave

Core observation: the LP's portfolio value is a **concave function** of the underlying price. By Jensen's inequality, the expected value of a concave function is less than the function at the expected value.

V2 LP value as function of price:

$$V_{\text{LP}}(P) = 2L\sqrt{P} = 2\sqrt{k\,P}$$

This is $\propto \sqrt{P}$ – a concave function.

HODL value:

$$V_{\text{HODL}}(P) = x_0 P + y_0$$

This is linear in P.

IL = concave – linear ≤ 0**:**

$$V_{\text{LP}}(P) - V_{\text{HODL}}(P) \leq 0 \quad \forall\, P$$

with equality only at $P = P_0$ (initial deposit price). The LP always loses relative to HODL.

V2 Impermanent Loss: Exact Formula

Uniswap V2 Impermanent Loss

$$\boxed{\mathrm{IL}(r) = \frac{2\sqrt{r}}{1+r} - 1 \leq 0, \qquad r = \frac{P_T}{P_0}}$$

Values:

r	IL	r	IL
0.5	-5.7%	1.5	-2.0%
0.25	-20.0%	2.0	-5.7%
0.1	-42.5%	5.0	-25.5%
0.01	-80.2%	10.0	-42.5%

Properties:

- Symmetric in r and $1/r$ (same IL for 2x up or 2x down)
- $\mathrm{IL} \to -1$ as $r \to 0$ or $r \to \infty$ (total loss in extreme moves)
- Second-order effect: $\mathrm{IL}(r) \approx -\frac{1}{8}(\ln r)^2$ for $r \approx 1$

Proof of V2 IL Formula

Setup: LP deposits at price P_0 into pool with $k = x_0 y_0$ and $P_0 = y_0/x_0$.

At price $P_T = rP_0$: pool rebalances via arbitrage to:

$$x_T = \sqrt{k/P_T}, \qquad y_T = \sqrt{k \cdot P_T}$$

LP value: (pro-rata share, assume 100% ownership for simplicity)

$$V_{\text{LP}} = x_T P_T + y_T = 2\sqrt{k\,P_T}$$

HODL value:

$$V_{\text{HODL}} = x_0 P_T + y_0 = \sqrt{k/P_0} \cdot P_T + \sqrt{kP_0}$$

Ratio:

$$\frac{V_{\text{LP}}}{V_{\text{HODL}}} = \frac{2\sqrt{kP_T}}{\sqrt{k/P_0} \cdot P_T + \sqrt{kP_0}} = \frac{2\sqrt{r}}{r+1}$$

$$\mathrm{IL} = \frac{V_{\text{LP}}}{V_{\text{HODL}}} - 1 = \frac{2\sqrt{r}}{1+r} - 1 \qquad \blacksquare$$

V3 Impermanent Loss: Echenim, Gobet, Maurice (2025)

Problem: V2 IL formula does not apply to V3 (concentrated liquidity changes everything).
Key paper: [5] – rigorous IL framework for Uniswap V3.

V3 IL for position in $[p_a, p_b]$, price moves $P_0 \to P_T$

$$\mathrm{IL}_{\text{V3}} = \frac{V_{\text{LP}}^{\text{V3}}(P_T)}{V_{\text{HODL}}(P_T)} - 1$$

where $V_{\text{LP}}^{\text{V3}}(P)$ is piecewise:

$$V_{\text{LP}}^{\text{V3}}(P) = \begin{cases} L\left(\frac{1}{\sqrt{p_a}} - \frac{1}{\sqrt{p_b}}\right)P & P < p_a \\ L\left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{p_b}}\right)P + L(\sqrt{P} - \sqrt{p_a}) & P \in [p_a, p_b] \\ L(\sqrt{p_b} - \sqrt{p_a}) & P > p_b \end{cases}$$

Middle case simplifies to: $V_{\text{LP}}^{\text{V3}}(P) = L\left(2\sqrt{P} - \sqrt{p_a} - P/\sqrt{p_b}\right)$

IL as the Gamma of a Self-Financing Portfolio

Key insight from [5]: impermanent loss is the *negative gamma* of the LP position.

Apply Itô's lemma to $V_{\mathrm{LP}}(P_t)$:

$$dV_{\mathrm{LP}} = \underbrace{\frac{\partial V}{\partial P}}_{\Delta} \,\mathrm{d}P + \frac{1}{2} \underbrace{\frac{\partial^2 V}{\partial P^2}}_{\Gamma} \sigma^2 P^2 \,\mathrm{d}t$$

For V2: $V = 2L\sqrt{P}$, so:

$$\Gamma = \frac{\partial^2 V}{\partial P^2} = -\frac{L}{2P^{3/2}} < 0$$

The IL accumulation rate is:

$$\boxed{\frac{d(\mathrm{IL})}{dt} = \frac{1}{2}\Gamma\,\sigma^2 P^2 = -\frac{L\sigma^2\sqrt{P}}{4} < 0}$$

Interpretation:

- LP has **negative gamma** – loses from price movement in either direction
- IL accumulates at rate proportional to σ^2 – **variance exposure**
- Identical to the "theta decay" of a short straddle in options

V3 IL Amplification: Concentration Multiplier

V3 concentrates liquidity → amplifies both fees AND impermanent loss.

V3 Gamma (within $[p_a, p_b]$):

$$\Gamma_{\text{V3}} = -\frac{L}{2P^{3/2}} \quad \text{(same formula, but } L \text{ is much larger per dollar invested)}$$

IL amplification factor:

$$\frac{\mathrm{IL}_{\text{V3}}}{\mathrm{IL}_{\text{V2}}} \approx \frac{L_{\text{V3}}}{L_{\text{V2}}} = \text{capital efficiency ratio}$$

Example: $\pm 10\%$ range, efficiency $\approx 10\times$:

- V2 IL for 20% move: -0.6%
- V3 IL for same move: $\approx -6\%$ (per dollar invested)
- But V3 also earns $\approx 10\times$ more fees

Net P&L of V3 LP

$$\text{Net P\&L} = \underbrace{\text{Fee revenue}}_{\text{positive, } \propto L \cdot \text{volume}} - \underbrace{|\text{IL}|}_{\text{negative, } \propto L \cdot \sigma^2}$$

Both scale with L – the LP bets that volume/variance ratio is favorable.

LP Position $\approx$ Short Straddle

Options analogy:

- **Short straddle:** sell ATM call + ATM put
- Payoff: $V(P) = -|P - P_0|$ (V-shape, concave)
- Collect premium (theta) while price stays near strike
- Lose if price moves significantly

LP position:

- Payoff $\propto \sqrt{P}$ – concave
- Collect fees ("theta") while price stays near P_0
- Lose (IL) if price moves significantly
- **Negative gamma, positive theta**

Greek comparison:

Greek	**Short Straddle**	**LP**
Delta	≈ 0 (ATM)	≈ 0.5
Gamma	< 0	< 0
Theta	> 0 (premium)	> 0 (fees)
Vega	< 0	< 0

V3 narrows the range:

LP $\approx$ short **strangle** with strikes p_a, p_b. Narrower range → higher gamma → larger "premium" (fees) required to compensate.

Fee-Adjusted P&L: When Do LPs Break Even?

LP net return: fees collected minus impermanent loss.

Break-even condition (continuous-time approximation)

$$\underbrace{\gamma \cdot \frac{\text{Volume}}{L}}_{\text{Fee rate}} \geq \underbrace{\frac{\sigma^2}{4\sqrt{P}}}_{\text{IL rate per unit } L}$$

Rearranging: the LP breaks even when

$$\frac{\text{Volume}}{\sigma^2} \geq \frac{L}{4\gamma\sqrt{P}}$$

Implications:

- High volume / low volatility → LP profitable (“fee farming”)
- Low volume / high volatility → LP loses (“toxic flow”)
- Stablecoin pairs (low σ): LPs tend to profit
- ETH/USDC (high σ): depends on volume regime

Empirical finding [11]: median Uniswap V2 LP earns slightly positive returns, but distribution is highly skewed – many LPs lose

Connection to Options Theory: LP as Exotic Derivative

Formal equivalence (Cartea et al. 2025 [2]):

A V2 LP position with initial deposit at price P_0 is equivalent to:

1. Hold $\frac{1}{2}$ unit of the risky asset (delta exposure)
2. **Plus:** continuously sell an infinitesimal straddle at the running price
3. The fee income plays the role of the straddle premium

V3 LP position in $[p_a, p_b]$**:**

- Equivalent to a **bull spread** (long call at p_a, short call at p_b) – in the token numeraire
- LP payoff at expiry: $V(P_T) = L[\min(\sqrt{P_T}, \sqrt{p_b}) - \sqrt{p_a}]^+$
- This is a **capped power option** – path-dependent due to fee accrual

Research frontier

Can we price and hedge LP positions using standard derivatives theory? Dynamic hedging of IL with options on Deribit?

Empirical IL: Lehar & Parlour (JF 2025)

Key findings from [11]: 19 million Uniswap V2 transactions.

LP profitability:

- Median LP earns small positive return (fees > IL)
- Mean LP return is **negative** (skewed distribution)
- Top 10% of LPs earn >50% annualized
- Bottom 10% lose >30% annualized

Who provides liquidity?

- Sophisticated LPs: active rebalancing, narrow ranges
- Passive LPs: "set and forget" – often lose
- Institutional LPs entering via V3

Who trades against the pool?

- Retail: informed by momentum, often "toxic"
- Arbitrageurs: 30% of volume, highly profitable
- MEV bots: sandwich attacks extract LP value

Fee revenue decomposition:

- Uninformed flow: positive for LPs
- Informed flow (arbitrage): negative for LPs
- Net: depends on pool and time period

Outline

1. CFMM Framework
2. Impermanent Loss
3. DeFi Lending
4. Applications & Industry

DeFi Lending Protocols: How They Work

Major protocols: Aave, Compound, Morpho, Spark (MakerDAO).

Mechanism:

1. **Lenders** deposit tokens into a pool → earn interest
2. **Borrowers** post collateral, borrow from pool → pay interest
3. **Interest rate** adjusts algorithmically based on **utilization**
4. **Liquidation** if collateral value drops below threshold

Key parameters:

- **Utilization:** $U = \frac{\text{Total Borrowed}}{\text{Total Supplied}}$
- **Collateral factor:** max borrow / collateral value (e.g., 80%)
- **Liquidation threshold:** triggers forced selling
- **Reserve factor:** protocol fee on interest

Scale (2025):

- Aave TVL: ∼\$20B
- Total DeFi lending: ∼\$40B
- Liquidation volume (2022 crash): >\$1B in 48h
- See [3]

Interest Rate Model: Piecewise Linear in Utilization

Standard DeFi interest rate curve (Aave V2/V3, Compound):

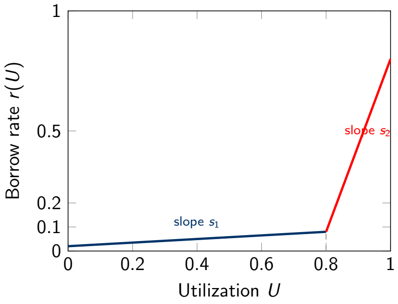

Piecewise linear rate:

$$r(U) = \begin{cases} r_0 + s_1 U & U \leq U^* \\ r_0 + s_1 U^* + s_2(U - U^*) & U > U^* \end{cases}$$

Parameters:

- r_0: base rate (∼0–2%)
- s_1: slope below kink (∼4–8%)
- U^*: optimal utilization (∼80%)
- s_2: slope above kink (∼300–700%)

Design goal: discourage $U > U^*$ (withdrawal risk for lenders). The steep slope above U^* acts as a "penalty rate."

Interest Rate Formulas: Full Specification

Borrow rate (Aave V3):

$$r_{\text{borrow}}(U) = \begin{cases} r_{\text{base}} + \dfrac{U}{U^*} \cdot r_{\text{slope1}} & \text{if } U \leq U^* \\ r_{\text{base}} + r_{\text{slope1}} + \dfrac{U - U^*}{1 - U^*} \cdot r_{\text{slope2}} & \text{if } U > U^* \end{cases}$$

Supply (lend) rate:

$$r_{\text{supply}}(U) = r_{\text{borrow}}(U) \cdot U \cdot (1 - \text{RF})$$

where RF = reserve factor (protocol revenue, typically 10–20%).

Typical parameters for USDC on Aave V3:

r_{base}	0%
r_{slope1}	4%
U^*	90%
r_{slope2}	60%
Reserve factor	10%

At $U = 80\%$: $r_{\text{borrow}} = 0 + (0.8/0.9) \cdot 4\% = 3.6\%$, $r_{\text{supply}} = 3.6\% \cdot 0.8 \cdot 0.9 = 2.6\%$.

Optimal Interest Rates: Stochastic Control

Key paper: [12] – "Optimal Risk-Aware Interest Rates for DeFi Lending."

Question: Can we derive the piecewise-linear rate from first principles as the *optimal* policy?

Setup

- Lending pool with total deposits D_t, total borrows B_t, utilization $U_t = B_t / D_t$
- Protocol sets borrow rate $r_t = r(U_t)$ as a function of utilization
- Borrowers arrive according to a point process; repay stochastically
- Lenders deposit/withdraw based on supply rate

State variable: utilization $U_t \in [0, 1]$.
Control: the interest rate function $r(\cdot)$.
Objective: maximize expected discounted lender wealth, penalizing insolvency risk.

Objective: Maximize Lender Wealth with Risk Penalties

Objective functional

$$\max_{r(\cdot)} \mathbb{E}\left[\int_0^T e^{-\beta t}\left[r(U_t)\,U_t\,(1-\mathrm{RF})\,D_t - \alpha\,\mathbf{1}_{\{U_t>\bar{U}\}}\,D_t\right]\mathrm{d}t\right]$$

Components:

- $r(U_t)\,U_t\,(1-\mathrm{RF})\,D_t$: interest income to lenders
- $\alpha\,\mathbf{1}_{\{U_t>\bar{U}\}}\,D_t$: penalty when utilization exceeds safety threshold $\bar{U}$
- β: discount rate
- T: planning horizon

Trade-off:

- Higher $r \rightarrow$ more interest income, but fewer borrowers (demand decreases)
- Lower $r \rightarrow$ more borrowers, but higher utilization $\rightarrow$ withdrawal risk
- Penalty term α captures the cost of "lender runs" (can't withdraw when $U \approx 1$)

Borrower Arrivals: Point Process with Intensity $\lambda(U)$

Borrower behavior:

- New borrows arrive as a point process N_t^+ with intensity $\lambda^+(r, U)$
- Repayments arrive as a point process N_t^- with intensity $\lambda^-(r, U)$
- Net utilization dynamics:

Utilization dynamics

$$dU_t = \underbrace{\lambda^+(r_t, U_t)\,\delta^+}_{\text{new borrows}}\,\mathrm{d}t - \underbrace{\lambda^-(r_t, U_t)\,\delta^-}_{\text{repayments}}\,\mathrm{d}t + \sigma_U\sqrt{U_t(1-U_t)}\,\mathrm{d}W_t$$

Intensity functions:

- **Linear:** $\lambda^+(r) = a - b\,r$ (demand decreases with rate)
- **Nonlinear:** $\lambda^+(r) = a\,e^{-b\,r}$ (exponential demand)
- Repayment rate: $\lambda^-(U) = c\,U$ (more borrows → more repayments)

The diffusion term $\sigma_U\sqrt{U(1-U)}$ ensures $U_t \in [0, 1]$ (Jacobi-type process).

HJB Equation for Optimal Rate

Value function: $V(t, U)$ = expected future discounted lender income starting from (t, U).

Hamilton-Jacobi-Bellman equation

$$\frac{\partial V}{\partial t} + \max_{r \geq 0} \left\{ \underbrace{\mu(r, U) \frac{\partial V}{\partial U}}_{\text{drift}} + \underbrace{\frac{1}{2}\sigma^2 U(1-U) \frac{\partial^2 V}{\partial U^2}}_{\text{diffusion}} + \underbrace{r\, U\, (1 - \text{RF}) - \alpha\, \mathbf{1}_{\{U > \bar{U}\}}}_{\text{running payoff}} \right\} = \beta V$$

where $\mu(r, U) = \lambda^+(r)\delta^+ - \lambda^-(U)\delta^-$ is the drift of utilization.

Optimal rate: first-order condition (pointwise maximization):

$$\frac{\partial}{\partial r}\left[\mu(r, U)\frac{\partial V}{\partial U} + r\, U\, (1 - \text{RF})\right] = 0$$

$$\Longrightarrow \qquad r^*(U) = \frac{U(1 - \text{RF}) + \frac{\partial \lambda^+}{\partial r}\delta^+ \cdot V_U}{-\frac{\partial^2 \lambda^+}{\partial r^2}\delta^+ \cdot V_U}$$

The optimal rate depends on the **shadow value of utilization** $V_U = \partial V / \partial U$.

Key Result: Optimal Rate is Bilinear under Linear Intensity

Theorem (arXiv 2502.19862, [ ])

Under **linear borrower intensity** $\lambda^+(r) = a - br$ with $a, b > 0$:
The optimal interest rate $r^*(U)$ is a **piecewise linear (bilinear) function** of utilization U, with a kink at a critical utilization U^*.

Explicit form:

$$r^*(U) = \begin{cases} r_0^* + s_1^* U & U \leq U^* \\ r_0^* + s_1^* U^* + s_2^*(U - U^*) & U > U^* \end{cases}$$

where r_0^*, s_1^*, s_2^*, U^* are determined by the model parameters $(a, b, \alpha, \beta, \sigma_U)$.

Remarkable implication

The piecewise-linear rate used by Aave/Compound is not an ad hoc design – it **emerges as the optimal policy** from a rational stochastic control problem!

Nonlinear Intensity: Monte Carlo + Deep Learning

When $\lambda^+(r)$ is nonlinear (e.g., exponential demand), the HJB does not admit a closed-form solution.

Numerical approaches [12]:

1. **Finite differences:** discretize (t, U) grid, solve HJB backward in time
   - Works for 1D state space
   - Curse of dimensionality for multi-pool problems
2. **Deep BSDE method:** represent $V(t, U)$ as a neural network
   - Forward simulate utilization paths
   - Backward propagate value via BSDE: $dV = [\text{driver}]\,\mathrm{d}t + Z\,\mathrm{d}W$
   - Train network to minimize terminal condition error
3. **Reinforcement learning:** model-free approach
   - Agent learns $r^*(U)$ by interacting with simulated pool environment
   - See [9] for RL in financial decision-making

Result: nonlinear optimal rates are “smoother” versions of the piecewise-linear form – the kink becomes a smooth transition.

Connection to Traditional Banking

Feature	Traditional Banking	DeFi Lending
Credit assessment	Manual / credit score	None (over-collateralized)
Collateral ratio	Often ¡ 100% (unsecured)	$>$ 110–150% (always secured)
Interest rate	Central bank + spread	Algorithmic (utilization-based)
Maturity	Fixed term	Open-ended (instant)
Liquidation	Bankruptcy process	Automatic (on-chain, instant)
Transparency	Opaque balance sheets	Full on-chain transparency
Regulation	Heavy (Basel III)	Minimal (evolving)
Deposit insurance	Government guarantee	None (smart contract risk)

Key insight: DeFi lending solves the **adverse selection** problem via over-collateralization, but creates **capital inefficiency**. The interest rate curve is a mechanism to manage utilization risk, analogous to reserve requirements in TradFi.

See [3] for a BIS perspective on why DeFi lending attracts borrowers.

Applications & Industry

Empirical evidence, risk management, industry practice

Lehar & Parlour (JF 2025): 19M Uniswap Transactions

Paper: “Decentralized Exchange: The Uniswap AMM” [11].

Data:

- 19 million transactions on Uniswap V2 and V3
- Complete on-chain data: every swap, mint, burn
- Period: May 2020 – Dec 2023

Methodology:

- Track individual LP positions from mint to burn
- Decompose LP returns into: fee revenue, IL, gas costs
- Classify traders: retail, arbitrageurs, MEV bots
- Compare V2 passive vs. V3 active LP strategies

Contribution: first comprehensive empirical study of AMM economics at scale, published in the *Journal of Finance*.

Key Findings: LP Returns, Trader Behavior, Fee Revenue

From [11]:

LP returns:

- Median LP: slightly positive (fees $>$ IL)
- Mean LP: **negative** (heavy left tail)
- Key driver: **adverse selection** – informed traders extract value from LPs
- V3 LPs: higher returns *if* actively managed; worse if passive

Fee revenue:

- Total fees: \$2–5B per year
- Fee/IL ratio varies by pool
- Stablecoin pools: consistently profitable
- ETH/USDC: marginal

Trader classification:

- Retail: 50% of trades, 20% of volume
- Arbitrageurs: 10% of trades, 30% of volume
- MEV bots: 5% of trades, 15% of volume (sandwich, backrun)
- Aggregators: 35% of trades, 35% of volume

Adverse selection:

- Arbitrage trades are “toxic” for LPs
- LPs lose on arb trades, profit on noise trades
- V3 concentratio makes adverse selection worse

V2 → V3 Transition: Concentrated Liquidity Effects

Findings on the V2/V3 comparison [11]:

Capital efficiency gains:

- V3 pools: 2–5× deeper at market price
- Lower slippage for traders
- LP capital requirements: 3–10× lower for same depth

LP management complexity:

- V3 requires active range management
- Gas costs for rebalancing: significant on Ethereum L1
- "Just-in-time (JIT) liquidity": sophisticated LPs add/remove liquidity within a single block

Impact on IL:

- V3 IL per dollar is higher (concentration effect)
- But V3 fee revenue per dollar is also higher
- Net effect: depends on range width and vol regime

Market structure shift:

- V3 favors professional LPs
- Passive retail LPs migrating to V2 forks or yield aggregators
- Ecosystem: LP management protocols (Gamma, Arrakis)

Aave V2 Empirics: BIS Working Paper 1183

Paper: Cornelli & Gambacorta (2024), "Why DeFi Lending?" [3].

Key findings:

- DeFi borrowers: mostly crypto-native (not accessing TradFi alternatives)
- Borrowing motives: leverage, tax optimization, yield farming
- Utilization: mean $\approx 60\%$, rarely exceeds optimal (U^*)
- Interest rate response: borrowers are rate-sensitive above U^*

Liquidation cascades:

- May 2022 (Luna): \$1B+ liquidated in 48h
- Liquidation penalty: 5–10% of collateral
- Cascade: liquidation → selling → price drop → more liquidations
- Systemic risk: [6]

Leverage in DeFi:

- Recursive borrowing: borrow, re-deposit, re-borrow
- Effective leverage: 2–5×
- See [8]

Gauntlet Network: Agent-Based Simulation for DeFi Risk

Industry practice: DeFi protocols use agent-based models (ABM) to stress-test parameters before governance votes.

Gauntlet's approach:

1. **Agent types:** borrowers (heterogeneous risk profiles), lenders, liquidators, arbitrageurs
2. **Price dynamics:** historical bootstrap or parametric (jump-diffusion)
3. **Protocol simulation:** exact smart contract logic replicated off-chain
4. **Outputs:** VaR, expected shortfall, liquidation cascades, protocol insolvency probability

Applications:

- Recommending collateral factors for new assets
- Stress testing interest rate curve parameters ($r_{\text{slope1}}, r_{\text{slope2}}, U^*$)
- Evaluating impact of oracle latency on liquidation efficiency
- Gauntlet manages risk for Aave, Compound, MakerDAO

Connection to our theory: ABM validates the stochastic control framework – optimal rates from HJB match empirically observed "good" parameter choices.

Interest Rate Curves: Governance Parameter Decisions

Who sets the parameters? Protocol governance (token holders vote).

Decision process:

1. Risk team (e.g., Gauntlet) proposes parameter change
2. Community discussion (Governance Forum)
3. On-chain vote (requires quorum)
4. Smart contract update (timelock delay)

Typical parameter updates:

- Increase U^* from 80% to 90% (after stability evidence)
- Reduce r_{slope2} to attract borrowers
- Add new collateral type with conservative LTV

Economic tensions:

- **Lenders** want: high rates, low risk
- **Borrowers** want: low rates, high LTV
- **Protocol** wants: high utilization (revenue), low insolvency risk
- **Token holders** want: TVL growth + protocol fees

Open research:

- Mechanism design: incentive-compatible parameter selection
- Dynamic adjustment: adaptive U^* based on market conditions
- Connection to [12]: implement optimal control on-chain?

Summary and Required Reading

Today we covered:

1. **CFMM Framework:** axiomatic theory (Angeris et al.), canonical trading functions
2. **Uniswap V2/V3:** constant product, concentrated liquidity, virtual reserves
3. **Impermanent Loss:** V2 formula, V3 amplification, IL as negative gamma
4. **LP $\approx$ short straddle:** connection to options theory
5. **DeFi Lending:** piecewise-linear rates, stochastic control (HJB), optimal bilinear rates
6. **Applications:** Lehar-Parlour (19M txns), BIS empirics, Gauntlet risk management

Required reading:

- [11] – Uniswap empirics (Journal of Finance)
- [1] – CFMM axiomatic framework (EC 2024)
- [5] – V3 impermanent loss (SIAM JFM)
- [12] – optimal lending rates (arXiv)

Day 3 preview: Microstructure & MEV – order flow, sandwich attacks, transaction fee mechanisms.

References I

[1] 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.

[2] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. "DeFi and Automated Market Makers: Execution and Speculation". In: *Journal of Economic Dynamics and Control* (2025).

[3] Giulio Cornelli and Leonardo Gambacorta. *Why DeFi Lending?* BIS Working Paper 1183. Bank for International Settlements, 2024.

[4] Marco Di Maggio. *Blockchain, Crypto and DeFi*. Wiley, 2024.

[5] Maxime Echenim, Emmanuel Gobet, and Anne-Claire Maurice. "Uniswap v3: Impermanent Loss Modeling". In: *SIAM Journal on Financial Mathematics* (2025).

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

[7] Campbell R. Harvey, Ashwin Ramachandran, and Joey Santoro. *DeFi and the Future of Finance*. Wiley, 2021.

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

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

[10] Albert S. Kyle. "Continuous Auctions and Insider Trading". In: *Econometrica* 53.6 (1985), pp. 1315–1335.

[11] Alfred Lehar and Christine A. Parlour. "Decentralized Exchange: The Uniswap Automated Market Maker". In: *Journal of Finance* 80.1 (2025), pp. 321–374.

[12] *Optimal Risk-Aware Interest Rates for DeFi Lending*. 2025. arXiv: `2502.19862 [q-fin.MF]`.