Problem Set 2: DeFi Mathematics

PhD Seminar – Digital Finance Research

Instructors: Prof. J. Osterrieder

Spring 2026

Due date: Two weeks from distribution.
Total points: 100.
Submission: PDF via course platform. Include all derivations and code (appendix or linked repository).
References: [6], [4], [1], [7], [3].

Problem 1: Impermanent Loss in Uniswap V3 [40 points]

In Uniswap V3, a liquidity provider (LP) deposits liquidity concentrated in the price range $[p_a, p_b]$ where $0 < p_a < p_b$. Let p_0 denote the price at deposit time and p the current price of asset X in terms of asset Y (i.e., $p = Y/X$).

Recall that a Uniswap V3 position in range $[p_a, p_b]$ with liquidity L holds reserves:

$$x(p) = L\left(\frac{1}{\sqrt{\hat{p}}} - \frac{1}{\sqrt{p_b}}\right), \tag{1}$$

$$y(p) = L\left(\sqrt{\hat{p}} - \sqrt{p_a}\right), \tag{2}$$

where $\hat{p} = \min(\max(p, p_a), p_b)$ is the price clamped to the active range.

(a) **(20 pts)** *Impermanent loss derivation.*

Define the value of the LP position at price p:

$$V_{\text{LP}}(p) = x(p) \cdot p + y(p).$$

Define the hold ("HODL") value as the value if the LP had simply held the initial reserves:

$$V_{\text{HODL}}(p) = x(p_0) \cdot p + y(p_0).$$

The impermanent loss is

$$\text{IL}(p) = \frac{V_{\text{LP}}(p)}{V_{\text{HODL}}(p)} - 1.$$

- (i) Compute $V_{\text{LP}}(p)$ and $V_{\text{HODL}}(p)$ explicitly for the three regimes: $p \le p_a$, $p_a < p < p_b$, and $p \ge p_b$.
- (ii) Derive a closed-form expression for $\text{IL}(p)$ when $p \in (p_a, p_b)$ and $p_0 \in (p_a, p_b)$.
- (iii) Let $\rho = p/p_0$ be the price ratio. Rewrite IL as a function of ρ, p_a/p_0, and p_b/p_0.

(iv) Show that for the full-range limit $p_a \to 0$, $p_b \to \infty$, your formula reduces to the classical Uniswap V2 impermanent loss:

$$\mathrm{IL}_{\mathrm{V2}}(\rho) = \frac{2\sqrt{\rho}}{1+\rho} - 1.$$

(b) **(10 pts)** *Numerical analysis.*

Fix $p_0 = 3{,}000$ (ETH/USDC).

(i) Plot $\mathrm{IL}(\rho)$ for three range widths: $[2500, 3500]$ (narrow), $[2000, 4000]$ (medium), $[1000, 6000]$ (wide), and the full-range V2 case.

(ii) For each range, compute the breakeven fee income: if the pool generates fee APY ϕ, find the annualized IL at which $\phi + \mathrm{IL}_{\mathrm{annual}} = 0$. Assume p follows a GBM with $\sigma = 80\%$ and compute $\mathbb{E}[\mathrm{IL}]$ via Monte Carlo (10,000 paths, 1-year horizon).

(c) **(10 pts)** *Connection to options.*

(i) Show that an LP position in Uniswap V3 with range $[p_a, p_b]$ can be replicated (up to a constant) by a portfolio of European-style power options (or covered calls/puts). Specifically, show that

$$V_{\mathrm{LP}}(p) = L\left(2\sqrt{\hat{p}} - \frac{\hat{p}}{\sqrt{p_b}} - \sqrt{p_a}\right)$$

and interpret each term.

(ii) Discuss: how does this options analogy inform hedging strategies for LPs? Cite [4].

Solution Sketch

(a)(i) For $p \in (p_a, p_b)$:

$$V_{\mathrm{LP}}(p) = L\left(\frac{1}{\sqrt{p}} - \frac{1}{\sqrt{p_b}}\right)p + L\left(\sqrt{p} - \sqrt{p_a}\right) = L\left(2\sqrt{p} - \frac{p}{\sqrt{p_b}} - \sqrt{p_a}\right).$$

For $p \le p_a$: $\hat{p} = p_a$, so $x = L(1/\sqrt{p_a} - 1/\sqrt{p_b})$, $y = 0$, and $V_{\mathrm{LP}} = L(1/\sqrt{p_a} - 1/\sqrt{p_b})\,p$.
For $p \ge p_b$: $\hat{p} = p_b$, so $x = 0$, $y = L(\sqrt{p_b} - \sqrt{p_a})$, and $V_{\mathrm{LP}} = L(\sqrt{p_b} - \sqrt{p_a})$.
(a)(ii) With both $p, p_0 \in (p_a, p_b)$:

$$\mathrm{IL}(p) = \frac{2\sqrt{p} - p/\sqrt{p_b} - \sqrt{p_a}}{\left(1/\sqrt{p_0} - 1/\sqrt{p_b}\right)p + \sqrt{p_0} - \sqrt{p_a}} - 1.$$

(a)(iii) Substituting $p = \rho p_0$, $\alpha = p_a/p_0$, $\beta = p_b/p_0$:

$$\mathrm{IL}(\rho) = \frac{2\sqrt{\rho} - \rho/\sqrt{\beta} - \sqrt{\alpha}}{\left(1 - 1/\sqrt{\beta}\right)\rho + 1 - \sqrt{\alpha}} - 1.$$

(a)(iv) As $\alpha \to 0$, $\beta \to \infty$: numerator $\to 2\sqrt{\rho}$, denominator $\to \rho + 1$, hence IL $\to 2\sqrt{\rho}/(1+\rho) - 1$.
(b) GBM simulation: $\ln(p_T/p_0) \sim \mathcal{N}((-\sigma^2/2)T, \sigma^2 T)$. With $\sigma = 0.8$, the expected annual IL is significant (∼5–15% for narrow ranges), motivating high fee APYs for profitability.
(c) The term $2L\sqrt{\hat{p}}$ behaves like a power perpetual (square-root payoff). The term $-Lp/\sqrt{p_b}$ is a short forward (capped), and $-L\sqrt{p_a}$ is a constant. This decomposition implies that LP positions have convexity exposure similar to short straddles, supporting hedging via options markets.

Problem 2: MEV Extraction as a Game [30 points]

Consider a simplified model of Maximal Extractable Value (MEV) extraction in a blockchain with Proposer–Builder Separation (PBS).

Setup. Two searchers (S_1 and S_2) compete to extract an arbitrage opportunity worth $V > 0$ (e.g., a DEX–CEX price discrepancy). Each searcher i chooses a bid $b_i \in [0, V]$ to submit to a single block builder. The builder includes the transaction with the highest bid.

(a) **(10 pts)** *Complete information game.*

Suppose both searchers observe V and simultaneously choose bids.

- (i) Write the payoff matrix for a discretized version where $b_i \in \{0, V/4, V/2, 3V/4, V\}$ (ties broken uniformly).
- (ii) Find all pure-strategy Nash equilibria.
- (iii) Find the mixed-strategy Nash equilibrium. Show that in the continuous limit ($b_i \in [0, V]$), the unique symmetric equilibrium has bids uniformly distributed on $[0, V]$, yielding expected profit $V/4$ per searcher and expected builder revenue $2V/3$.

(b) **(10 pts)** *Asymmetric information.*

Now suppose S_1 has latency advantage: with probability α, S_1 observes S_2's bid before choosing.

- (i) Model this as a Bayesian game. Write the expected payoffs.
- (ii) Find the equilibrium bidding strategies as a function of α.
- (iii) Show that as $\alpha \to 1$ (full latency advantage), S_1's expected profit approaches $V/2$ and S_2's approaches 0.

(c) **(10 pts)** *Implications for PBS design.*

- (i) The builder's revenue in part (a) is $2V/3$. Compare this to the first-price sealed-bid auction revenue (with two bidders, uniform values on $[0, V]$) of $V/3$. Explain the discrepancy.
- (ii) Discuss how the MEV auction design relates to the analysis in [2] and [7]. What are the welfare implications of MEV extraction for ordinary users?
- (iii) Propose a mechanism modification that would reduce searcher profits and increase builder (or user) surplus. Analyze whether it is incentive-compatible.

Solution Sketch

(a)(i) Payoff matrix for S_1 (row) vs. S_2 (column), entry = (S_1 profit, S_2 profit):

	0	$V/4$	$V/2$	$3V/4$	V
0	$(V/2, V/2)$	$(0, 3V/4)$	$(0, V/2)$	$(0, V/4)$	$(0, 0)$
$V/4$	$(3V/4, 0)$	$(3V/8, 3V/8)$	$(0, V/2)$	$(0, V/4)$	$(0, 0)$
$V/2$	$(V/2, 0)$	$(V/2, 0)$	$(V/4, V/4)$	$(0, V/4)$	$(0, 0)$
$3V/4$	$(V/4, 0)$	$(V/4, 0)$	$(V/4, 0)$	$(V/8, V/8)$	$(0, 0)$
V	$(0, 0)$	$(0, 0)$	$(0, 0)$	$(0, 0)$	$(0, 0)$

(a)(ii) Pure NE: (V, V) is the only pure NE (both bidding V, zero profit) – but also $(3V/4, V)$ and $(V, 3V/4)$ are NE since neither can profitably deviate. The game resembles an all-pay auction.

(a)(iii) In the continuous all-pay auction with two symmetric bidders and common value V, the unique symmetric equilibrium has CDF $F(b) = b/V$ on $[0, V]$. Expected profit per searcher: $\int_0^V (V-b)\cdot(b/V)\cdot(1/V)\,\mathrm{d}b = V/4$. (This uses the order statistics of two uniforms.) Builder revenue: $\mathbb{E}[\max(b_1, b_2)] = 2V/3$ by the expectation of the maximum of two Uniform$(0, V)$ draws.

(b) With probability α, S_1 observes b_2 and bids $b_2+\epsilon$. In the "no observation" case, both play the symmetric equilibrium. In equilibrium, S_2 shades bids downward anticipating S_1's advantage. As $\alpha \to 1$, S_1 always wins and pays just above S_2's bid, extracting nearly all surplus.

(c) The discrepancy arises because MEV is a common-value setting (both know V), not a private-value auction. Oz, Sui, Tirole (2025) show MEV extraction creates allocative inefficiency. Possible mechanisms: commit-reveal bidding, order flow auctions returning MEV to users, or encrypted mempools (threshold encryption).

Problem 3: VPIN Estimation for Crypto Markets [30 points]

The Volume-Synchronized Probability of Informed Trading (VPIN) metric, adapted from the PIN framework, measures toxicity of order flow. You will implement VPIN on synthetic crypto trade data.

Background. Partition total volume into *buckets* of equal size $\bar{V}$. Within each bucket n, classify trades as buyer- or seller-initiated and compute:

$$\text{VPIN}_n = \frac{1}{\ell} \sum_{j=n-\ell+1}^{n} \frac{|V_j^B - V_j^S|}{\bar{V}},$$

where V_j^B and V_j^S are buyer and seller volumes in bucket j, and ℓ is the lookback window (number of buckets).

(a) **(10 pts)** *Algorithm description.*

Write detailed pseudocode for the VPIN computation:

  (i) Trade classification: apply the tick rule (compare consecutive trade prices) to label each trade as buy or sell.

  (ii) Volume bucketing: accumulate trades into buckets of exactly $\bar{V}$ units. Handle partial fills at bucket boundaries (i.e., a single trade can span two buckets).

  (iii) VPIN computation: rolling average over ℓ buckets.

  (iv) Discuss: why is volume-time superior to clock-time for high-frequency crypto data? Cite [3].

(b) **(10 pts)** *Implementation on synthetic data.*

Generate 1,000 synthetic trades as follows:

- Initial price $p_0 = 50{,}000$ (BTC/USD).
- Each trade: $\Delta p_k \sim \mathcal{N}(0, 10^2)$ (price change), volume $v_k \sim \text{Exp}(\mu = 0.5)$ BTC.
- With probability 0.3, inject an "informed" block: 5 consecutive trades with $\Delta p_k \sim \mathcal{N}(\pm 50, 5^2)$ (directional pressure) and $v_k \sim \text{Exp}(\mu = 2.0)$.

Parameters: $\bar{V} = 10$ BTC, $\ell = 20$ buckets.

  (i) Implement the VPIN algorithm in Python.

  (ii) Plot the VPIN time series (in volume-time). Overlay the informed trading episodes.

(iii) Does VPIN spike during informed episodes? Report the mean VPIN during informed vs. uninformed periods and run a Welch t-test.

(c) **(10 pts)** *Interpretation and extensions.*

(i) What VPIN threshold would you use as an early warning for toxic flow? Discuss Type I vs. Type II errors.

(ii) Compare VPIN to the Kyle lambda (λ_K) from [5]. Under what conditions do they measure the same thing?

(iii) How would VPIN behave in a DEX (e.g., Uniswap) vs. a centralized exchange? Discuss the role of MEV-related toxic flow in DEX settings, referencing [7].

Solution Sketch

(a) Pseudocode:

```
function COMPUTE_VPIN(trades, V_bar, ell):
    # Step 1: Classify trades via tick rule
    for k = 1 to N:
        if p[k] > p[k-1]: side[k] = BUY
        elif p[k] < p[k-1]: side[k] = SELL
        else: side[k] = side[k-1]  # same as previous

    # Step 2: Bucket trades by volume
    buckets = [], current_buy = 0, current_sell = 0, current_vol = 0
    for k = 1 to N:
        remaining = v[k]
        while remaining > 0:
            space = V_bar - current_vol
            fill = min(remaining, space)
            if side[k] == BUY: current_buy += fill
            else: current_sell += fill
            current_vol += fill
            remaining -= fill
            if current_vol >= V_bar:
                buckets.append((current_buy, current_sell))
                current_buy = 0, current_sell = 0, current_vol = 0

    # Step 3: Compute rolling VPIN
    vpin = []
    for n = ell to len(buckets):
        s = sum(|B_j - S_j| / V_bar for j in [n-ell+1..n])
        vpin.append(s / ell)
    return vpin
```

Volume-time is superior because it normalizes for the extreme variation in crypto trading activity across time-of-day and news events.
(b) Typical results: VPIN $\approx$ 0.3–0.4 during normal periods, spiking to 0.6–0.8 during informed blocks. Welch t-test: expect $p < 0.001$ for the difference in means.
(c)(i) A threshold of 0.5–0.6 balances Type I/II errors; the exact value should be calibrated to the specific market.
(c)(ii) VPIN and Kyle lambda both measure adverse selection. Under Kyle's model with a single informed trader and Uniform prior, $\lambda_K = \sigma_v/(2\sigma_u)$ where σ_v is fundamen-

tal volatility and σ_u is noise trader volume. VPIN is a non-parametric estimator that converges to λ_K-like quantities under specific structural assumptions.

(c)(iii) In DEX settings, all transactions are public on-chain. MEV searchers front-run or sandwich ordinary traders, creating toxic flow that would register as elevated VPIN. However, the lack of a traditional order book complicates direct application.

References

[1] Guillermo Angeris, Tarun Chitra, Theo Diamandis, Alex Evans, and Gerry Kulkarni. The geometry of constant function market makers. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC 2024)*, 2024.

[2] Soroush Bahrani, Pranav Garimidi, and Tim Roughgarden. Transaction fee mechanism design in a post-MEV world. In *Proceedings of the 6th Conference on Advances in Financial Technologies (AFT 2024)*, 2024.

[3] David Easley, Maureen O'Hara, Liyan Yang, and Zhuo Zhang. Microstructure and market dynamics in crypto, 2024. SSRN 4814346.

[4] Maxime Echenim, Emmanuel Gobet, and Anne-Claire Maurice. Uniswap v3: Impermanent loss modeling. *SIAM Journal on Financial Mathematics*, 2025.

[5] Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

[6] Alfred Lehar and Christine A. Parlour. Decentralized exchange: The Uniswap automated market maker. *Journal of Finance*, 80(1):321–374, 2025.

[7] Baran Oz, Simon Sui, and Jean Tirole. Maximal extractable value and allocative inefficiency. *Journal of Financial Economics*, 2025.