SoK: Impermanent Loss, An Unavoidable Fee or a Controlled Phenomenon?

Arad Kotzer
Technion

Ori Rottenstreich
Technion

***Abstract*—Decentralized exchanges built on Automated Market Maker (AMM) protocols have become a cornerstone of Decentralized Finance (DeFi), offering token swaps without conventional order-book matching. However, supplying liquidity to these AMM pools exposes participants to distinctive market risks, most notably *impermanent loss* (IL) - the potential underperformance of a liquidity-provider portfolio relative to simply holding the underlying tokens. This paper presents a comprehensive overview of IL, unifying its main theoretical models, empirical evidence, and mitigation strategies. Our survey spans constant-function and concentrated liquidity market makers, synthesizes findings from leading DeFi protocols, and reviews mitigation methods that include both protocol-level adaptations and financial-engineering approaches. Across these perspectives, we highlight recurring trade-offs, cost, complexity, and security. Finally, we outline open questions for managing IL's systemic effects, keeping decentralized liquidity provision both profitable for participants and sustainable for the broader ecosystem.**

I. Introduction

The Decentralized Finance (DeFi) ecosystem comprises various financial services that run as blockchain applications [1]–[6]. DeFi enables participants to earn yield by providing services (e.g., lending assets for interest or supplying liquidity for fees). The core strategies revolve around lending money for interest and providing liquidity to token exchanges for transaction fees. They also include generating actions on DeFi platforms to earn engagement rewards. Estimates of the DeFi token market cap range from approximately 90 billion USD [7] to 103 billion [8], as of December 2025.

Many DeFi applications rely on the provision of liquidity. Liquidity providers are typically compensated through interest payments, trading fees, or other incentive mechanisms. However, supplying liquidity entails exposure to various forms of risk. In lending protocols, for instance, lenders face the possibility of borrower default or delayed repayment [9]–[11]. A central challenge specific to AMM-based exchanges is *impermanent loss* (IL). *IL is the shortfall between the value of an LP's pool position and the value of simply holding the same tokens (HODL).* IL arises when a token's market price changes, leading to a rebalancing of the pool's asset composition. This rebalancing reduces the relative value of the liquidity provider's position. The magnitude of IL increases with the extent of price divergence. This loss is not realized unless assets are withdrawn, hence it is "impermanent", and it can theoretically disappear if token prices return to their original state. To emphasize the severity of IL, Loesch et al. [12] analyze 17 popular Uniswap v2 pools covering 43% of the locked total value of the market, between May and September 2021. They show that the total fees earned since inception were 199M USD, while simply holding the tokens would have yielded 260M USD —an underperformance of roughly 61M USD.

Example. We illustrate IL using real ETH values from May 2025. Table I compares the outcomes of providing liquidity versus a HODL strategy. An LP initially deposits 10000 USDC and 5 ETH, with ETH priced at 2000 USDC, for a total value of 20000 USDC. When ETH appreciates to 2500 USDC, arbitrage trading rebalances the pool so that the LP holds 4.4721 ETH and 11180.3 USDC, worth 22360.68 USDC. By contrast, under HODL, the assets would be worth 22500 USDC. The difference of about 139.3 USDC constitutes the IL, despite the LP earning an absolute profit compared to the initial investment.

This fundamental risk raises a central question: Is impermanent loss an unavoidable consequence of on-chain market making, or can it be actively mitigated?

This Systematization of Knowledge synthesizes theory and empirical evidence, suggesting a nuanced answer: IL is a *structural* property of AMM mechanics. However, its realized impact can be materially reduced through a spectrum of strategies. These include protocol-level design choices and financial overlays, each with distinct trade-offs in cost, complexity, and security. **Thus, IL is unavoidable in principle, but can be controlled in practice within these trade-offs.**

Contributions. This paper systematizes the research on IL in DeFi and makes the following key contributions:

1) **Unified View of IL Formalizations.** We unify the primary mathematical models of IL. This shows how the choice of performance benchmark critically impacts its

Table I
Comparison of liquidity provision and HODL strategy for ETH–USDC pool, May 2025.

	Liquidity initial	**HODL value**	**Liquidity value**
USDC	10000	10000	11180.3
ETH	5	5	4.4721
USDC price	1	1	1
ETH price	2000	2500	2500
Total value (USD)	20000	22500	22360.68
IL (USD)	22500 – 22360.68 ≈ 139.32		

measurement and interpretation (Section III).

2) **Critical Synthesis of Empirical Findings.** We synthesize and reconcile conflicting empirical findings on LP profitability. We trace the reported divergences to key methodological factors, including fee-reinvestment models, market-volatility regimes, and the choice of asset pools (Section IV).
3) **Taxonomy of Mitigation Strategies.** We propose a novel taxonomy of IL mitigation strategies based on their layer of intervention. *Protocol-level mechanisms* modify an AMM's core design, while *financial-overlay instruments* use derivatives for hedging. This framework systematically exposes the distinct trade-offs in cost, complexity, and security for each category (Sections V–VI).
4) **System-Wide Perspective and Open Directions.** We analyze IL from a holistic, system-wide perspective. This analysis maps the economic interdependencies between LPs, traders, arbitrageurs, and adjacent DeFi protocols. Based on this analysis, we outline a forward-looking research agenda to address critical open questions (Section VII).

The remainder of the paper is structured as follows. In Section II, we review core concepts of AMMs, liquidity provision, and IL. Section III formalizes IL: we present canonical CPMM formulas and multi-asset extensions, incorporate fees and invariance principles, extend to alternative AMM designs, and discuss benchmark choices. Section IV synthesizes empirical evidence on LP profitability across major deployments. Section V reviews protocol-level approaches, including novel pricing rules, oracle-assisted AMMs, order-book-inspired designs, adaptive fees, and weighted pools. Section VI examines mitigation via financial engineering, including option-based hedges and volatility-linked protection claims. Finally, Section VII discusses the results, solution trade-offs, and open questions for future research. Table II summarizes the surveyed works by subject and protocol. We use standard abbreviations (e.g., IL: Impermanent Loss, IG: Impermanent Gain, CPMM: Constant Product Market Maker, LVR: Loss-Versus-Rebalancing), which are defined throughout the text.

II. Background

A. Automated Market Makers

Automated Market Makers (AMMs) [48] are decentralized exchange protocols that enable coin exchange using *liquidity pools*. Liquidity pools are smart contracts wherein Liquidity Providers (LPs) lock assets to facilitate trading. The most common AMM type is the *constant function market maker* (CFMM), where a trading function $f(x_0, x_1)$ imposes an invariant on the token reserves x_0, x_1. The most prominent example is the constant-product rule $x_0 \cdot x_1 = K$, used globally by Uniswap v2 and locally (within specific price bounds) by Uniswap v3. Crucially, this invariant dictates that the instantaneous exchange rate is determined by the ratio of the reserves (x_1/x_0). Generalizations include weighted geometric means (Balancer) and stableswap hybrids that blend constant product and sum (Curve) [49].

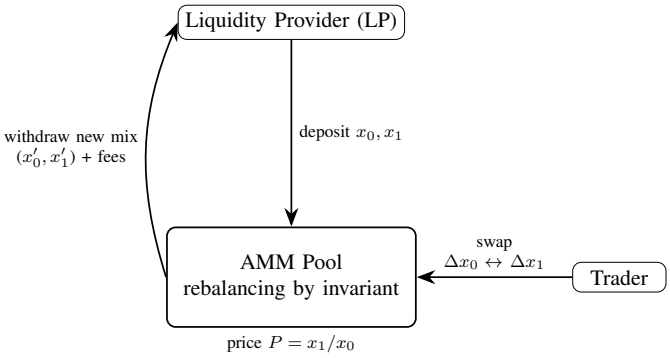

Figure 1. Liquidity provision illustration of a constant-product AMM

AMMs effectively act as *algorithmic price oracles*. Whenever the pool price deviates from external markets, arbitrage restores alignment [50]. Arbitrage thus serves as both a stabilizing mechanism and a source of value extraction [51]. This arbitrage-based alignment sustains on-chain liquidity and determines value transfers among participants. Traders may pay *slippage*, defined as execution at an average price worse than the pre-trade marginal price. Conversely, LPs may incur IL when their pool position underperforms an equivalent HODL benchmark. Prior work links these two notions. Specifically, in constant-product AMMs, trader slippage corresponds to the mirror image of LP impermanent gain [42]. Moreover, adversarial ordering (e.g., front-running and sandwiching) can extract value under *maximal extractable value* (MEV) [52]–[56].

B. Liquidity Provision

Liquidity provision is the process by which users deposit tokens into an AMM pool, enabling traders to swap assets directly against pooled reserves. In return, LPs receive a proportional share of the pool's swap fees [10], [11], [57]. For example, in a Uniswap v2 ETH/USDC pool, LPs must deposit token *amounts* of equal *value*. Traders then exchange one asset for the other via the constant product invariant of the pool. This mechanism underpins decentralized exchanges by providing continuous on-chain liquidity without a traditional order book. Figure 1 illustrates the liquidity provision flow. First, an LP deposits amounts (x_0, x_1) of assets (c_0, c_1) into the pool. When a trader swaps, the pool balances change along the pricing curve, updating the implied price. Finally, upon withdrawal, the LP receives her new proportional mix of tokens plus the fees accumulated by the position.

a) Fee microstructure and concentrated liquidity: AMM protocols differ in how fees interact with liquidity. In Uniswap v2, swap fees are *reinvested into the pool reserves* on each trade, so the product $x_0 \cdot x_1 = K$ grows. In Uniswap v3, fees are added to the *separate fee counters* of each position and do not change the active liquidity L or the reserves until the LP explicitly claims or reinvests them. Moreover, v3 allows

Subject	Relevant References
IL Definitions & Formalizations	[13] Tiruviluamala et al. (canonical CPMM & n-asset IL; PLI/ERLI), [14] Aigner & Dhaliwal (IL vs. initial portfolio; v3 range-dependence), [15] Boueri (volatility-driven IL in geometric-mean AMMs), [16] Roman et al. (fee-adjusted IL/IG; constant-K vs. fee-sensitive), [17]–[19] (early/intro IL discussions).
Alternative Benchmarks & Bounds	[20] Milionis et al. (early LVR framing), [21] Milionis et al. (LVR: slippage/arbitrage cost; hedged-LP; variance link), [22] Cohen et al. (hedging/inefficiency bound for fair fees), [23] Tangri et al. (RV/FARV; fee-adjusted relative value).
Empirical Analyses of IL	
Uniswap v2 & forks	[24] Heimbach et al. (v2, pool categories & IL), [15] Boueri (147 v2 pools, median IL), [25] Hafner & Dietl (agent-based on v2 WETH/USDC), [16] Roman et al. (Uniswap/Sushiswap; fee-adjusted IG prevalence).
Uniswap v3	[12] Loesch et al. (v3, minimal/out-of-range/actual IL vs. fees), [26] Heimbach et al. (v3 positions; risk/return & range effects), [23] Tangri et al. (v3 RV/FARV; passive underperformance), [16] Roman et al. (v3; IG prevalence with fees).
Balancer v2	[23] Tangri et al. (trade-level profitability; active LPs), [16] Roman et al. (fee-adjusted outcomes), [27] Lebedeva et al. (adaptive-fee simulations).
Curve	[16] Roman et al. (fee-adjusted outcomes; smaller IG share).
Hedging IL with Financial Derivatives	[28] Deng et al. (static replication with vanilla options), [29] Khakhar & Chen (LP–PNL delta-style hedging), [30] González et al. (CPMM long-strangle replication), [31], [32] Fukasawa et al. (weighted variance swaps for IL; vol-derivatives background), [33] Lipton et al. (IL protection claim; funded vs. borrowed).
AMM Design Comparisons & Innovations	[34] Kim et al. (CPMM/CSMM/CCMM/LS-LMSR; IL vs. IG), [13] Tiruviluamala et al. (general CFMM framework), [35] Lee & Kim (transaction-level pricing; batching/ordering), [36] Tran et al. (OBMM: square-root impact & reversals; IL elimination conditions), [27] Lebedeva et al. (adaptive fee schedules), [37] Nadkarni et al. (oracle-free *adaptive curves*, Glosten–Milgrom zero-profit; Kalman filtering; v4 hooks + ZK), [38] Baggiani et al. (optimal dynamic fees; two-regime policy; inventory-linear, oracle-tracking fees), [39] DODO (oracle-based PMM), [40] Chen et al. (MPMM, multi-token PMM), [41] Pattanaik (weighted pools & IL trade-offs).
Ecosystem Interactions & Indirect Drivers of IL	[42] Labadie (IL–slippage equivalence; value transfer), [43] Wang et al. (leveraged LPing; propagation via lending), [44], [45] Aave documentation (liquidations), [46], [47] liquidation/oracle-driven shocks that impact AMM prices.

Table II
CLASSIFICATION OF IL-RELATED PAPERS BY SUBJECT

LPs to choose a price band $[p_-, p_+]$ and provide liquidity only inside it. When $p_{\text{pool}} \in [p_-, p_+]$, the position behaves locally like a CPMM with "virtual" reserves determined by $[p_-, p_+]$ and its liquidity L, and it earns fees. If price exits this band, the position converts to a single asset (c_0 if $p_{\text{pool}} < p_-$, c_1 if $p_{\text{pool}} > p_+$) and *accrues no fees* until price re-enters or the LP rebalances. Figure 2 illustrates the liquidity distribution in Uniswap v2 and v3.

C. Impermanent Loss

Impermanent loss (IL) is a common risk AMM LPs face. When token prices change, the invariant rebalances reserves. This results in a different asset mix than initially deposited, creating a valuation gap regardless of whether the LP withdraws immediately. Consequently, the LP's position can underperform a simple buy-and-hold (*HODL*) portfolio of the same initial assets [17]–[19], [23], [42]. We formally define IL as

Definition 1 (Impermanent Loss)**.** *IL is the relative underperformance of the position of an LP in an AMM compared to the HODL strategy.*

IL arises from unequal changes in the value of the tokens inside the liquidity pool. Arbitrageurs align AMM prices with external market prices, and in doing so, crystallize IL for LPs. While impermanent loss is typically framed as a shortfall relative to a HODL benchmark, recent work [58] interprets it more generally as a form of *price impact* and value transfer. Under this view, IL can be seen as a signed quantity whose effect depends on how trades move AMM prices relative to the external market: Trades that move prices toward the market tend to generate losses for LPs, whereas trades that move prices away from the market may generate gains. This perspective helps explain why several studies have identified, in particular CFMM settings, the existence of *Impermanent Gain* (IG), namely the positive relative performance of a liquidity-

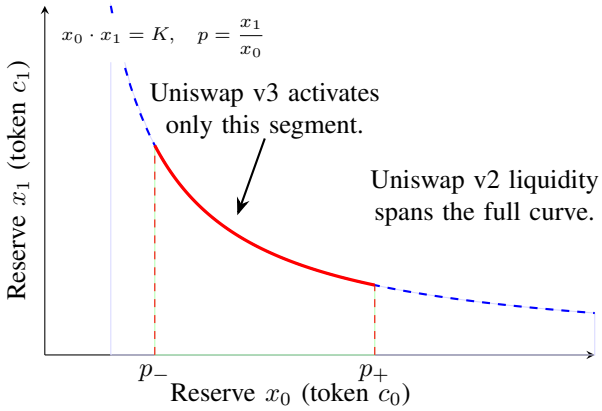

Figure 2. Liquidity distribution in Uniswap v2 (full range) vs. Uniswap v3 (concentrated in $[p_-, p_+]$). The curve shows a constant-product contour $x_0 \cdot x_1 = K$, and the highlighted segment marks the price band where v3 liquidity is active.

provider position compared to the HODL strategy [16], [34].

III. IL Formalization

We now survey the principal mathematical formulations of IL proposed in the literature. These definitions both establish how IL is measured and provide extensions to different AMM designs. This section is organized into four parts. Canonical IL formulas, generalizations and fee adjusted metrics, different AMM designs, and alternative benchmarks and perspectives. We conclude the section with a short synthesis that shows how these expressions fit into a single value based template.

A. Canonical Formulas of IL

The simplest and most widely used *expression* for IL arises in constant-product market makers (CPMMs) such as Uniswap v2. Consider a two-asset pool with reserves (x_0, x_1) and instantaneous price $p = x_1/x_0$, satisfying the invariant $x_0 \cdot x_1 = K$. We focus on the canonical $50{:}50$ *(equal-value)* case, meaning the LP provides equal value for each asset at the reference time. Hence, the corresponding benchmark is a HODL position with the same initial dollar split.

Under continuous rebalancing and in the absence of fees, Tiruviluamala et al. [13] show that IL depends only on the price ratio $R = p'/p$, where p is the initial price and p' is the *final* price at the end of the holding period:

$$\mathrm{IL}(R) = \frac{2\sqrt{R}}{1+R} - 1, \qquad R = \frac{p'}{p}. \tag{1}$$

This canonical IL is always non-positive, i.e., $\mathrm{IL}(R) \leq 0$ for all $R > 0$. Negative values indicate that the LP position underperforms the HODL benchmark, while $\mathrm{IL}(R) = 0$ occurs only when $R = 1$ (no price change). In the limit of large price divergence ($R \to 0$ or $R \to \infty$), $\mathrm{IL}(R) \to -1$, corresponding to a maximal relative loss of 100%.

Moreover, IL is symmetric in *log-returns*. Letting $\ell = \ln R$, we have $\mathrm{IL}(R) = \mathrm{IL}(1/R)$, i.e., the loss depends only on $|\ell|$ rather than on the direction of the price move. Figure 3 illustrates $\mathrm{IL}(R)$ as a function of the price change. This

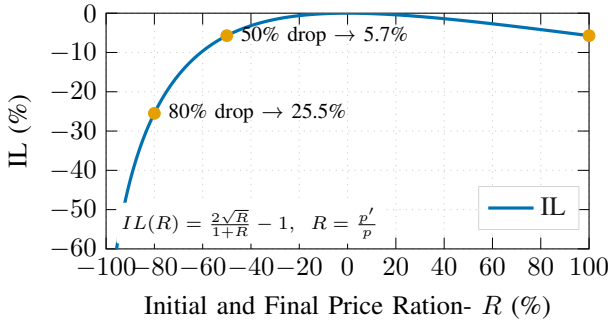

Figure 3. IL vs. Price Change, for a $50{:}50$ CPMM pool (no fees).

canonical IL is always non-positive, i.e., $\mathrm{IL}(R) \leq 0$ for all $R > 0$. Negative values indicate that the LP position underperforms the HODL benchmark, while $\mathrm{IL}(R) = 0$ occurs only when $R = 1$ (no price change). In the limit of large price divergence ($R \to 0$ or $R \to \infty$), $\mathrm{IL}(R) \to -1$, corresponding to a maximal relative loss of 100%.

The paper later extends this definition to n-asset (geometric-mean / constant-mix) pools. Let $R_i = p_i'/p_i$ denote the relative price change of asset i against a chosen numeraire (so one can take $R_0 = 1$). For equal weights, IL can be written as

$$\mathrm{IL} = \frac{\left(\prod_{i=1}^n R_i\right)^{1/n}}{\frac{1}{n}\sum_{i=1}^n R_i} - 1 \leq 0, \tag{2}$$

For $n = 1$, the expression is trivially 0, since there is no second asset to rebalance against. For $n = 2$ with $(R_1, R_2) = (R, 1)$, (2) reduces to (1).

The paper introduces two invariance principles: *Price Level Independence* (PLI) and *Exchange Rate Level Independence* (ERLI). PLI means IL is unchanged by a common rescaling of all prices, so it depends only on exchange rates rather than absolute price levels. ERLI strengthens this by requiring that IL depends only on *changes* in those relative prices (i.e., on the ratios R_i), not on their initial levels. Among common AMM designs, geometric-mean mechanisms satisfy both properties, which makes them particularly analytically tractable.

a) Volatility-driven IL: Boueri [15] further shows that in geometric mean AMMs, IL grows systematically with volatility. Under the assumption that arbitrage eliminates profit opportunities, IL relative to a constant-mix portfolio of weights w (representing the fixed proportion of the pool's total value assigned to each asset) is defined as

$$IL(t) = \exp\left(\frac{1}{2}\int_0^t \left(\langle w|\Sigma w\rangle - \langle w|\sigma^2\rangle\right) ds\right), \tag{3}$$

where $\langle a \mid b\rangle \triangleq a^\top b$ denotes the standard inner product. In particular, $\langle w \mid \Sigma w\rangle = w^\top \Sigma w$ is the portfolio variance under weights w, while $\langle w \mid \sigma^2\rangle = \sum_i w_i \sigma_i^2$ is the weighted average of individual variances. Since their difference is always non-positive, portfolios in AMMs systematically underperform

their constant-mix benchmarks. This behavior is analogous to *Gamma Selling* [1], where higher volatility exacerbates IL.

B. Generalizations and Fee-Adjusted Metrics

Recent work extends IL definitions to cover broader CFMM families and incorporate trading fees.

a) Relative Value and FARV: Tangri et al. [23] present a general mathematical framework for IL in any N-asset CFMM, including pools with optional concentrated liquidity. The authors introduce the notion of *Relative Value* (RV) and *Fee-Adjusted Relative Value* (FARV). RV measures LP performance against HODL, while FARV includes fee income. The paper proves that LP profitability requires $V_i^\lambda \geq 1$ where V_i^λ denotes the FARV of asset c_i. The authors additionally extend these equations to concentrated liquidity positions.

b) Fee-driven IL and IG: Roman et al. [16] formalize how transaction fees affect the shape and even the sign of IL, combining theory with Ethereum data. In AMMs such as Uniswap v2 [59], Balancer [60], and Curve [61], fees are reinvested into the pool. Hence, they increase the reserves and change the liquidity factor K. In Uniswap v3 [62], on the other hand, fees are distributed separately and K remains constant. Hence, the authors claim some existing non-Uniswap v3 analyses may not be accurate.

The authors define IL after a single transaction as the difference between the deposit value and the HODL allocation:

$$V_{\text{deposit}} = p(x_0 + \Delta x_0, K)(x_0 + \Delta x_0) + (x_1 - \Delta x_1), \quad (4)$$

$$V_{\text{HODL}} = p(x_0 + \Delta x_0, K) \cdot x_0 + x_1, \quad (5)$$

with absolute and relative IL given by

$$IL_a = V_{\text{deposit}} - V_{\text{HODL}} = p(x_0 + \Delta x_0, K)\Delta x_0 - \Delta x_1, \quad (6)$$

$$IL_r = \frac{V_{\text{deposit}}}{V_{\text{HODL}}} - 1. \quad (7)$$

where IL_a and IL_r denote the absolute and relative impermanent loss, respectively. Here, IL_a measures the absolute value gap, while IL_r expresses the loss (or gain) relative to the HODL benchmark, enabling comparison across pools and trade sizes. This makes IL directly comparable across pools and trade sizes. Negative values indicate impermanent losses, while positive values imply impermanent gains. The paper additionally derives closed-form expressions for these fee-adjusted metrics for Uniswap and Balancer (whereas Curve requires numerical solutions). Crucially, with non-zero fees, IG arises for small trades (typically when the pool fee rate satisfies $\alpha < 0.2\%$ of reserves).

C. Range-dependent IL in Uniswap v3

Unlike v2's full-range liquidity, Uniswap v3 allows LPs to choose a price interval $[p_-, p_+]$ and supply liquidity only within that band. Based on these mechanics, Loesch et al. [12] define three IL notions. These describe different states of the LP position along a price path (in-range, out-of-range, and realized), rather than additive components of loss. A numerical illustration appears in Table III.

- **Minimal (in-range) IL**: The unavoidable loss while the price remains within $[p_-, p_+]$;
- **Out-of-range IL**: Additional loss once the position becomes inactive (in principle avoidable by withdrawing / rebalancing);
- **Actual IL**: The realized total (including time out of range).

Heimbach et al. [26] extend IL formalization to concentrated liquidity pools.

$$IL_{\text{v3}}(p, p'; p_-, p_+) = \frac{V_{\text{LP}}(p'; p_-, p_+)}{V_{\text{HODL}}(p')} - 1. \quad (8)$$

where p and p' are the initial and final market prices (relative to the base coin), V_{LP} is the terminal value of a position active only on $[p_-, p_+]$, and V_{HODL} is the corresponding HODL value. Closed-form expressions reveal the connection between the definitions: IL_{v3} functions as a leverage-amplified version of the canonical IL. As the range narrows, the effective leverage increases, making the loss magnitude significantly larger than in v2 for the same price deviation.

D. Different AMM Designs

Kim et al. [34] extend the IL analysis beyond the pools of constant products to cover constant sum (CSMM), constant circle (CCMM), and liquidity-sensitive logarithmic market scoring rule (LS-LMSR) invariants. This comparison highlights how IL and IG manifest differently across AMM families.

The invariants are (Figure 4 illustrates example functions of AMMs):

- Constant Product (CPMM): $x_0 \cdot x_1 = K$
- Constant Sum (CSMM): $x_0 + x_1 = K$
- Constant Circle (CCMM): $(x_0 - a)^2 + (x_1 - b)^2 = K$
- Liquidity-Sensitive LMSR (LS-LMSR): $C = b \cdot \ln\left(e^{x_0/b} + e^{x_1/b}\right) = K$

The paper evaluates the value of the gap between pool and HODL portfolio values (IL). Explicit formulas are derived for CPMM and CSMM, while LS-LMSR and CCMM yield more complex or implicit expressions (summarized in Table IV).

To interpret the following equations, we define the notation as follows: x_1 is the initial reserve of the first token, x_1' is the updated reserve of that token after rebalancing, and k is

Table III
ILLUSTRATIVE EXAMPLE OF MINIMAL, OUT-OF-RANGE, AND ACTUAL IL FOR THE ETH/USDC POOL. ALL VALUES ARE EXPRESSED IN USDC UNITS.

	HODL value	LP value	IL
Initial deposit	20,000	20,000	0
Within range (Minimal IL)	22,500	22,360.7	139.3
Out-of-range state (inactive position)	25,000	15,000	10,000
Actual IL (total)	25,000	15,000	10,139.3

[1] In traditional finance, providing liquidity in a CFMM is mathematically equivalent to a *Short Gamma* position. Sellers of such positions profit when price volatility is low but face accelerating losses (convexity) during large, directional price swings.

Table IV
SUMMARY OF IL AND IMPERMANENT GAIN (IG) ACROSS DIFFERENT AMM TYPES, BASED ON KIM ET AL. [34].

AMM Type	Invariant	Closed-form IL?	Outcome
CPMM	$x_0 \cdot x_1 = K$	Yes (quadratic)	IL only
CSMM	$x_0 + x_1 = K$	Yes (quadratic)	IL + IG
CCMM	$(x_0 - a)^2 + (x_1 - b)^2 = K$	No (implicit form)	IL + IG
LS-LMSR	$C = b \cdot \ln\left(e^{x_0/b} + e^{x_1/b}\right) = K$	No (high-order)	IL + IG

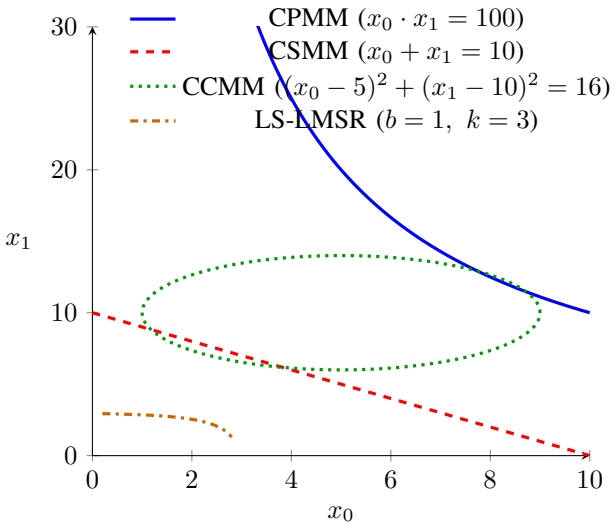

Figure 4. Comparison of AMM invariants with parameters: CPMM ($k = 100$), CSMM ($k = 10$), CCMM (center $(a, b) = (5, 10)$, radius $r = 4$), and LS–LMSR ($b = 1$, $k = 3$).

the constant invariant of the pool. The IL is expressed as a function of the updated reserve x_1':

$$IL_{\text{CPMM}}(x_1') = \frac{1}{{x_1'}^2}\left[-\frac{K}{x_1}{x_1'}^2 + 2K\,x_1' - Kx_1\right], \tag{9}$$

$$IL_{\text{CSMM}}(x_1') = \frac{1}{x_1'}\Big[-2{x_1'}^2 + (K + 2x_1)x_1' - Kx_1\Big], \tag{10}$$

$$IL_{\text{LS-LMSR}}(x_1') \ \text{ no closed form (high-order function)}, \tag{11}$$

$$IL_{\text{CCMM}}(x_1') \ \text{ implicit via } x_1' = b + \sqrt{K - (x_0 - a)^2}. \tag{12}$$

E. Alternative Benchmarks and Perspectives

Beyond the canonical HODL comparison, a growing body of work proposes alternative benchmarks and related metrics. They seek to isolate particular risk factors and offer more intuitive views of liquidity-provider performance.

a) LVR benchmark: Milionis et al. [20], [21] introduce *Liquidity Value at Risk (LVR)* as a benchmark that captures a different aspect of AMM risk than canonical IL. Intuitively, IL measures the underperformance relative to a passive HODL strategy, whereas LVR compares the AMM outcome to a continuously rebalanced portfolio. Thus, LVR isolates the cost of adverse selection arising from arbitrage. The two notions are related but use different benchmarks: IL compares the AMM position to a passive HODL strategy, whereas LVR compares it to a continuously rebalanced portfolio. As a result, they can yield different values under the same price path.

Let p_t denote the external market price of the risky asset at time t. Let $V(p)$ be the CFMM pool-value function evaluated at price p, so that the realized pool value at time T is $V_T \equiv V(p_T)$.

The paper compares V_T to a *continuously rebalanced* reference portfolio (a constant-mix strategy) whose value is

$$V_T^{\text{RB}} = V_0 + \int_0^T x^*(p_s)\,dp_s, \qquad T \geq 0.$$

Here,

- $V(p)$ is the CFMM pool-value function at price p (hence $V_T = V(p_T)$),
- $x^*(p)$ is the risky-asset position of the rebalanced benchmark as a function of the market price p (i.e., the holdings required to maintain the target constant-mix).

The *LVR* is defined as the shortfall of the actual pool value relative to this rebalanced benchmark:

$$\text{LVR}_T \ = \ V_T^{\text{RB}} - V_T.$$

Equivalently, it measures how much value the AMM loses compared with an ideal portfolio that is continuously rebalanced to maintain the same target mix.

The authors show that

$$\text{LVR}_T \ = \ \int_0^T \ell(p_s)\,ds, \qquad \ell(p) \ \triangleq \ -\frac{\sigma_p^2\,p^2}{2}\,V''(p) \ \geq \ 0, \tag{13}$$

where σ_p is the instantaneous volatility of p_t. Thus, LVR is predictable, non-decreasing, and directly quantifies the adverse-selection cost from arbitrage, framing LPs as effectively short volatility.

Hedging bound. Complementary to LVR, Cohen et al. [22] propose a hedging view that treats an LP's exposure as a perpetual option and derives an arbitrage-free *upper bound* on fee income needed to offset IL fairly. Multi-agent simulations with fees in the 0–4% range typically find fee flow below this bound, with losses that increase as volatility rises.

b) Redefinition of IL: Aigner and Dhaliwal [14] reformulate IL by evaluating it relative to the *initial* portfolio value rather than the final one. For Uniswap v2 they derive

$$IL(R) = \sqrt{R} - \frac{R+1}{2} = -\tfrac{1}{2}(\sqrt{R}-1)^2 \tag{14}$$

This definition yields smaller losses for downward moves (e.g., -4.29 vs. -5.72 for a 50% price drop) but larger losses for upward moves, aligning more closely with investor intuition. Extending to Uniswap v3, they show IL is *range-dependent*, requiring quadratic solutions for reserves at the price bounds. Numerical examples confirm that tighter ranges amplify both risk and asymmetry: A $\pm 20\%$ move results in about -0.5 IL in v2, but nearly -4.8 in a v3 position constrained to $[80\%, 120\%]$.

F. Synthesis via a Unified Value Template

The subsections above followed the notation and fee accounting conventions of their respective source papers. We close by rewriting these definitions in a single notation, so that the expressions in Section III can be read as special cases of one underlying comparison.

Let $\mathbf{p} = (p_1, \ldots, p_n)$ be the vector of token prices at the end of the holding period, and let $\mathbf{p}^{\text{in}}$ be the entry price vector. Let θ collect AMM and position parameters such as the invariant family, weights, or a concentrated liquidity range. Denote by $x^{\text{LP}}(\mathbf{p};\theta) \in \mathbb{R}^n_+$ the terminal LP holdings after arbitrage has aligned the pool position with market prices $\mathbf{p}$. The LP marked-to-market value is

$$V^{\text{LP}}(\mathbf{p};\theta) \triangleq \langle \mathbf{p},\, x^{\text{LP}}(\mathbf{p};\theta)\rangle. \tag{15}$$

A benchmark is defined by entry-time holdings $x^{\text{B}}(\mathbf{p}^{\text{in}};\theta)$. Its terminal value at prices $\mathbf{p}$ is

$$V^{\text{B}}(\mathbf{p};\mathbf{p}^{\text{in}},\theta) \triangleq \langle \mathbf{p},\, x^{\text{B}}(\mathbf{p}^{\text{in}};\theta)\rangle. \tag{16}$$

We then define Relative Value (RV) and impermanent loss (IL) by

$$\begin{aligned} \text{RV}(\mathbf{p};\mathbf{p}^{\text{in}},\theta) &\triangleq \frac{V^{\text{LP}}(\mathbf{p};\theta)}{V^{\text{B}}(\mathbf{p};\mathbf{p}^{\text{in}},\theta)}, \\ \text{IL}(\mathbf{p};\mathbf{p}^{\text{in}},\theta) &\triangleq \text{RV}(\mathbf{p};\mathbf{p}^{\text{in}},\theta) - 1. \end{aligned} \tag{17}$$

This view organizes the literature around three choices. First, the holdings map x^{LP} is determined by the AMM mechanism and any position constraints, such as full-range versus concentrated liquidity. Second, the benchmark holdings x^{B} encode the reference strategy, which is HODL in the canonical definition but may be a continuously rebalanced portfolio in alternatives such as LVR. Third, fee handling specifies whether fees are compounded into the LP position or accrued separately and added when evaluating performance.

a) Mapping surveyed formulas to this template: Table V summarizes how the main notions discussed above fit into (17). The CPMM definition in (1) is obtained by using the constant product holdings map with the HODL benchmark formed at entry. The geometric mean expression in (2) follows the same construction for n assets under a constant-mix mechanism. Fee-adjusted metrics such as RV and FARV correspond to evaluating (17) under an explicit fee convention, which determines how fee income enters V^{LP}. Finally, alternative benchmarks such as LVR primarily modify the benchmark component. Although LVR is usually written as an absolute value gap rather than a ratio, it is still based on comparing the pool value function against a rebalanced reference portfolio.

IV. Empirical Analysis

Empirical studies of IL evaluate how LPs perform, based on real data. While some analyses conclude that IL often exceeds fee income, others find that dynamic effects can offset or even reverse these losses. In this section, we synthesize the main empirical findings and their implications for LP profitability.

A. Evaluation of IL

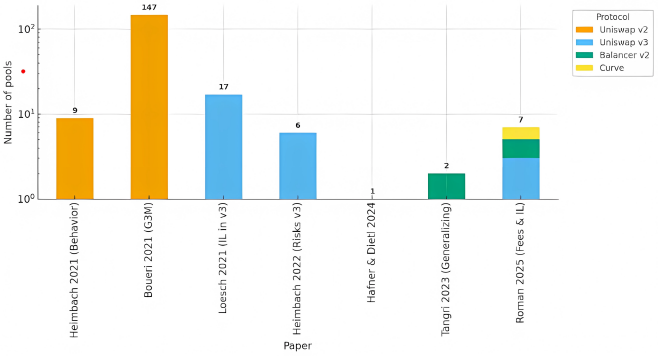

Figure 5. IL analysis: Pools by paper and protocol.

Several key studies have quantified the real-world impact of IL across different protocols and market conditions, often arriving at conflicting conclusions.

Analyzing over 29,000 Uniswap v2 pools, *Heimbach et al.* [24] demonstrated that LP profitability is decisively shaped by asset type. The paper focuses on nine pools in the context of IL: Three stablecoin pools, three of "normal" (standard non-pegged asset) tokens, and three "exotic" ones. While stablecoin pools provided consistent low-risk returns, "exotic" pools suffered from severe IL that fees could not offset. This highlights the vulnerability of LPs in highly volatile markets. *Boueri et al* [15] verify that impermanent loss is inherent in Uniswap v2–style pools.

Research on Uniswap v3 provides further evidence of IL. *Loesch et al.* [12] analyzed 17 major pools and found that IL consistently exceeded fee income, challenging the assumption that concentrated liquidity inherently solves the problem. Their findings also suggest that active management of liquidity ranges did not significantly improve outcomes. Similarly, *Heimbach et al.* [26] concluded that Uniswap v3's design transforms liquidity provision into a high-risk activity, favoring sophisticated players. Volatile pools exposed LPs to extreme IL when using narrow ranges, with fewer than 30% of positions being adequately compensated for the additional risk.

Other studies have focused on the methodologies used to measure IL, revealing that outcomes depend heavily on the

Table V
MAPPING FROM SURVEYED NOTIONS TO THE UNIFIED VIEW IN (17).

Study or notion	Mechanism in $x^{\mathrm{LP}}(\mathbf{p};\theta)$	Benchmark in x^{B}	Fee handling
CPMM IL (1)	Full range CPMM	HODL	None
Geometric mean (2)	Constant mix weights	HODL	Typically none
RV and FARV	CFMM or v3 position	HODL	By convention
Uniswap v3 IL (8)	Concentrated range	HODL	Often segregated
LVR	Same pool value function	Rebalanced constant mix	Optional

analytical framework. *Tangri et al.* [23] introduce a generalized framework (FARV). The paper finds that while passive LPs in Uniswap v3 and Balancer v2 tended to underperform the HODL strategy, a trade-level analysis revealed that the majority of individual swaps were profitable. This highlights a significant potential for active management.

Departing from static analyses, *Hafner & Dietl* [25] use an agent-based model. The paper argues that IL is better understood as a rare outcome of extreme and sustained price shocks (e.g., $> 75\%$ annual drops). The paper's dynamic model finds that arbitrage activity can benefit LPs, because frequent trades generate enough fee income to offset potential IL.

Similarly, *Roman et al.* [16] focus on the impact of reinvested fees. The paper shows that conventional constant-K models systematically overestimate IL. Their analysis of Uniswap, Balancer, and Curve shows that small trades often produce IG rather than losses. In most Uniswap and Balancer pools studied, more than 95% of swaps yielded IG, highlighting the need for fee-sensitive models to assess risk accurately. Table VI presents a comparison between the papers, and Figure 5 summarizes the pools used by each paper.

B. Reconciling Findings: When Do Fees Offset IL?

Section IV-A presents some conflicting results: while some studies find that IL consistently dominates fees [12], [26], others report that fee income frequently restores profitability [16], [25]. Synthesizing these conflicting reports reveals that divergences are typically driven by five primary methodological factors:

1) **Benchmark and Metric:** Reported outcomes depend critically on the performance benchmark. Comparing the pool value against a simple HODL strategy often shows losses in trending markets due to the pool's convex payoff structure [12], [20]. In contrast, metrics that generalize relative value or fee-adjusted returns demonstrate that LPs can remain profitable even under high terminal price divergence if trading volume is sufficient [23].
2) **Fees and Microstructure:** The treatment of fees significantly alters results. In AMMs where fees are reinvested into the reserves (e.g., Uniswap v2), the growth of the liquidity parameter K creates a "buffer" that can result in *Impermanent Gain* (IG) during periods of high volume and low volatility [16]. Conversely, models that assume a constant-K or segregate fees (e.g., Uniswap v3) report significantly higher exposure to IL because the core liquidity does not benefit from auto-compounding [24].
3) **Pool Selection and Market Regimes:** Profitability is decisively shaped by asset volatility. Studies finding severe underperformance often focus on "exotic" or highly volatile pairs during sustained directional trends [15], [24]. In contrast, studies reporting favorable outcomes often examine stablecoin pairs or high-liquidity pools characterized by mean-reverting price action, where fee accumulation has time to outpace price drift [25], [26].
4) **Strategy and Infrastructure:** The choice of liquidity range and the underlying chain's efficiency are crucial. Passive LPs in concentrated liquidity pools face "convexity risk"; narrow ranges amplify IL magnitude unless positions are actively rebalanced [12]. Furthermore, infrastructure-level factors such as block times impact arbitrage frequency; lower frictions and different block intervals on various networks may generate higher fee velocity for LPs compared to the high-friction environment of certain mainnets [21], [24].
5) **Liquidity timing (JIT):** *Just-in-time* (JIT) liquidity occurs when an LP adds liquidity immediately before a swap and removes it immediately after. This lets tactical LPs capture that swap's fees with minimal inventory risk, often at the expense of passive LPs. Empirically, JIT is concentrated in a few addresses and is relatively infrequent. However, when it does occur, it increases near-price depth and can improve taker execution (lower slippage), while reducing the fee share earned by HODL positions. Under standard assumptions, the private gain from a single JIT event is bounded by about 2α (one fee on entry and one on exit around the swap) [63], [64].

Rules of thumb.

(1) *Stable or tightly correlated pairs* with sustained volume tend to see fees offset IL under fee-reinvestment [15], [24].

(2) *Concentrated liquidity* increases convexity. Narrow ranges increase IL unless actively and correctly rebalanced. Fees often fail to compensate at the position level [12], [26].

(3) *Fee-aware accounting* and *trade-level views* can flip many small swaps to IG [16], [23]. However, this does not guarantee profitable positions once gas/MEV and miss-rebalancing are included.

(4) *Dynamic or agent-based* analyses suggest fees can dominate except under extreme sustained trends [25], and in

Table VI
EMPIRICAL IL RESULT COMPARISON.

Study	Benchmark	Fees	Pools	Period	Method	Outcome
Heimbach [24]	IL vs HODL	Reinvested	Stable / Normal / Exotic	2020–21	Ex-post on-chain	Severe IL in exotic
Boueri [15]	IL vs HODL	Reinvested	Uniswap v2	2020–21	Ex-post	Median IL ∼20%
Loesch [12]	Min/Actual IL	Separate	Uniswap v3	2021	Ex-post	IL > fees in most pools
Heimbach [26]	IL vs HODL	Separate	Stable / Volatile	2021–22	Model + data	Stable: rare losses, Volatile: common losses
Hafner–Dietl [25]	LP P&L	Reinvested	v2 WETH/USDC	2021	Agent-based	Profitable except extreme trends
Roman [16]	Fee-adjusted IL/IG	Reinvested	v2 / Balancer / Curve	2024	Ex-post	IG on >95% of swaps (small trades)
Tangri [23]	FARV / RV	Mixed	v3 / Balancer	2023	Pool + trade level	Passive underperforms, active helps

general high price volatility leads to IL [22].

V. MECHANISM DESIGN APPROACHES TO IL MITIGATION

IL is a central concern in on-chain market making. A large body of work aims to *redesign* AMMs - Adjusting pricing, fees, or liquidity allocation - to reduce IL while preserving market quality. We group mechanisms into four families.

A. Transaction-Level Pricing

Lee & Kim [35] distinguish the *execution price* of a trade (the average price implied by reserve changes) from the LP's *value gain* after a batch, valued at external prices. The gain is then evaluated as $G = p_1(x_1' - x_1) + p_0(x_0' - x_0)$, where $G > 0$ indicates IG. The execution price governs the path from (x_0, x_1) to (x_0', x_1'). By reordering and batching swaps (a combinatorial problem conjectured to be NP-hard by the authors), LPs can skew paths for higher G. Empirically, batching reduces IL, slippage, and gas, reframing part of MEV as an optimization lever rather than purely adversarial extraction.

B. Oracle-Based Solutions

DODO [39] is a DeFi protocol for token swaps and liquidity provision. Its Proactive Market Maker (PMM) departs from fixed constant-product formulas by using price oracles and a tunable curvature parameter K to align pool prices with external market prices. This design concentrates liquidity near the current price, leading to higher capital efficiency, lower slippage, and reduced IL compared to traditional CPMMs. Building on this idea, Chen et al. [40] introduce the Multi-Token Proactive Market Maker (MPMM), which extends PMM to handle multiple assets.

Using hourly price data for 14 tokens (April 2021–September 2023) and interpolated trading volumes from CoinGecko [8], the authors run large-scale simulations under real-world scenarios. Results show that MPMM consistently outperforms PMM, CPMM, and multi-token CPMMs in stable and declining markets, delivering higher efficiency and lower IL. However, in bull markets, MPMM adapts more slowly to rapid appreciation, and in crash scenarios, volatility propagates across all assets in the pool.

Beyond market-specific performance, oracle-based AMMs introduce distinct operational risks: **(i) Security risk:** vulnerability to oracle manipulation attacks, such as price-lag exploitation; **(ii) Latency risk:** the time delay between external market moves and on-chain updates, which creates "stale" prices that arbitrageurs can pick off; and **(iii) Liveness risk:** the protocol's dependence on a continuous, reliable external data feed to maintain the pool's internal price alignment.

C. Order-Book–Inspired AMMs

Tran et al. [36] propose an *Order-Book–Inspired AMM* (OBMM) that incorporates statistical features from centralized limit order books (LOBs) to mitigate IL. Unlike CPMMs, which exhibit asymmetric liquidity, quadratic price impact, and zero price reversal, the OBMM introduces three empirically grounded features: (i) Symmetric liquidity depth around the mid-price; (ii) A concave, square-root price impact rule; (iii) Partial price reversal, where only a fraction of a trade's price impact persists.

The paper demonstrates that these adjustments, particularly a square-root impact combined with a partial reversal factor ($\lambda \leq 2/3$), can eliminate IL per trade and potentially generate IG. Empirical validation using real-world and synthetic datasets confirms that the OBMM reduces IL and lowers slippage on large trades compared to CPMMs.

D. Adaptive Fee Mechanisms

Fixed fees expose LPs during high volatility. Lebedeva et al. [27] evaluate block-adaptive, deal-adaptive, and oracle-based rules that raise fees when adverse selection risk is high, redirecting value from informed arbitrage toward LPs. On historical paths (e.g., ETH/SHIB), block-adaptive improves markouts by $\approx 1.5\%$, while oracle-based reduces IL by $> 7\%$.

Another approach derives *oracle-free, adaptive bonding curves* that enforce a Glosten–Milgrom [65] zero-expected-profit condition using only trade history, explicitly targeting zero-profit conditions and LVR. Nadkarni et al. [37] characterize the optimal adaptive curve using a differential equation

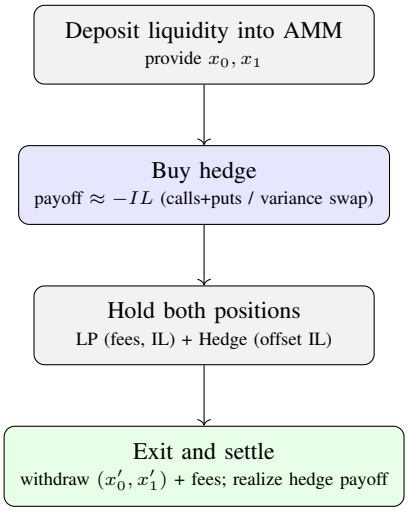

Figure 6. General hedge-based mitigation scheme: (1) Deposit liquidity, (2) Buy a hedge with payoff $\approx -IL$, (3) Hold both positions, (4) Redeem tokens and settle the hedge.

and implement it through Kalman filtering [66]. The paper reports lower arbitrage losses than static CFMMs.

E. Weighted Pool Designs

Pattanaik [41] analyzes weighted Balancer pools through a simplified two-asset framework. By exploiting exchange-rate level independence, IL is expressed as a function of the relative price change and its pool weight w_1:

$$L(R_1, w_1) = \frac{(R_1)^{w_1}}{R_1 \cdot w_1 + (1 - w_1)} - 1 \tag{18}$$

Graphical and sensitivity analysis reveal that evenly weighted pools (e.g., $w_1 = 0.5$) incur the largest IL under price divergence, while skewed pools (e.g., $w_1 = 0.8$) substantially reduce losses. However, uneven weights also increase slippage, reducing utilization and fee income. The study highlights a central trade-off: IL can be lowered by adjusting pool weights, but at the cost of liquidity efficiency.

VI. IL Mitigation Using Financial Derivatives

A complementary line of work hedges IL with *off-chain financial derivatives*. Two views dominate: (i) IL as an option-like payoff (replicate with calls/puts), and (ii) IL as *weighted variance* exposure. Figure 6 summarizes the operational flow (deposit, buy hedge, hold, settle).

A. Option-Like IL Formalization

Deng et al. [28] formalize IL in Uniswap v3 as a portfolio of vanilla options and derive static replication formulas that allow LPs to hedge with European options [67]. The key equation, *unit IL per unit liquidity* (UIL), is shown to be equivalent to a weighted basket of calls and puts over the liquidity interval,

$$\mathbb{E}[\mathrm{UIL}_R] = -\tfrac{1}{2} \int_{P_\ell}^{P_u} K^{-3/2} C(K)\, dK,$$

with an analogous expression for the lower bin. The proposed strategy requires buying a static option portfolio at pool entry and holding until withdrawal, thus avoiding rebalancing. Monte Carlo simulations [68] under a Heston model [69] show replication errors as small as 10^{-6}–10^{-5}. Tests with Deribit Bitcoin options [70] confirm that accurate hedges remain feasible even with only a few strikes. The study highlights that IL can be fully understood through classical option sensitivities and mitigated with one-time static hedges.

Khakhar et al. [29] shift the focus from IL to the *Liquidity Position PNL* (LP–PNL), defined as the actual return on the LP's position relative to initial investment:

$$\text{LP-PNL} = \frac{\text{Final value of pool assets}}{\text{Initial value of LP investment}} - 1.$$

For Uniswap v2, this admits a closed form in terminal price change, while for Uniswap v3, it depends on final token balances within the liquidity range. The authors propose to hedge LP–PNL by using a sparse static portfolio of vanilla options, tuned by least-squares regression across possible terminal prices. Empirical analysis using ETH and WBTC Deribit options shows that the combined portfolio largely cancels out the LP's sensitivity to price movements across wide ranges. This implies that LPs can retain fee income while remaining largely neutral to directional market risk.

Gonzalez et al. [30] further develop an option-based static hedge for CPMMs such as Uniswap and Balancer. They represent the pool's terminal value as $V_T = 2\sqrt{kp'}$. The paper then expresses IL as a nonlinear payoff function of p'. By decomposing this payoff into a portfolio of European calls and puts, the authors show that LP exposure can be replicated through a *long strangle* strategy. The paper proves that, under appropriate inequalities relating to option strikes, premiums, and pool returns, LPs can guarantee non-negative payoffs within a pre-specified price interval. A numerical illustration with real Ethereum options data from Deribit confirms the practical viability of this approach.

This supports two main conclusions. First, IL is ubiquitous in deployed AMMs. It arises as a structural consequence of continuous rebalancing under stochastic relative prices, appearing across diverse designs and major deployments. Second, fees can partially hedge IL but do not universally eliminate it. Net outcomes depend on the interaction between fee generation and divergence dynamics. Therefore, periods of high volume may still coincide with negative net performance.

Beyond LP profitability, IL also has broader ecosystem implications. Execution in CFMMs is mediated by price impact, leading to value transfers to arbitrageurs and slippage paid by traders. In constant-product settings, this slippage can be interpreted as the mirror image of LP impermanent gain. This relationship highlights a redistribution channel between traders and LPs [42]. Moreover, liquidity positions are often leveraged, for example, by borrowing against LP shares. In these cases, divergence-driven drawdowns can be amplified into liquidation risk. This dynamic potentially propagates stress across interconnected DeFi protocols.

B. Hedging IL using Protection Claims

Another line of work links IL to volatility derivatives and develops direct protection instruments. A central tool is the *weighted variance swap* [32], a contract that exchanges realized variance against a fixed strike, where the variance is computed with a weight function of the underlying price. By choosing the weight, one can target specific volatility exposures, mapping IL to a tradable payoff.

Fukasawa et al. [31], [32] show that, for a two-asset CPMM with relative price process p_T (token 1 in units of the base token) and an equal weight LP position at inception ($50:50$ by value), the *dollar value* of IL at time T (measured in units of the base token) can be written as

$$\mathrm{IL}(T) \;=\; x_1 \left(\sqrt{\frac{p_T}{p_0}} - 1 \right)^2 . \tag{19}$$

This is equivalent to shorting a weighted variance swap with a weight function $g(p) = \sqrt{p}$. Hence, IL is an exposure to realized volatility (not just a static value gap), because it depends on both the magnitude of price changes and the price level. Using replication arguments, the weighted variance swap can be hedged with strips of vanilla options. Close approximations are obtainable with standard variance or gamma swaps when weighted contracts are unavailable. Extensions include AMM fees (which reduce effective IL) and generalized AMM designs whose IL can be engineered to match classical volatility-swap payoffs.

Similarly, Lipton et al. [33] treat IL itself as a derivative payoff, an *IL protection claim*. For Uniswap v2 they derive closed forms (in base units), distinguishing funded vs. borrowed LP positions:

$$\mathrm{IL}^{\mathrm{funded}}(p_T) = \sqrt{\frac{p_T}{p_0}} - 1,$$

$$\mathrm{IL}^{\mathrm{borrowed}}(p_T) = -\frac{1}{2}\left(\sqrt{\frac{p_T}{p_0}} - 1 \right)^2 .$$

Define the protection payoff as $\Pi_T \equiv -\mathrm{IL}(p_T)$. This equals the LP's losses and can be valued with standard derivative-pricing machinery. For Uniswap v3, with active range $[p_\ell, p_u]$, IL admits a decomposition into a linear term, a square-root payoff over $[p_\ell, p_u]$, and option payoffs at the range boundaries. This enables static replication with out-of-the-money options and dynamic hedging under stochastic volatility. Hedge costs increase as the liquidity range narrows.

VII. Discussion

IL is well-recognized in AMMs. The literature reviewed in this paper indicates that IL is not a rare phenomenon; rather, it is a systematic outcome of AMM mechanics that appears across diverse pools, protocols, and market conditions. Transaction fees can sometimes offset these losses, but rarely do so consistently, leaving LPs structurally exposed. Meanwhile, mitigation research spans protocol redesigns, oracle-based pricing, and derivative overlays. While these approaches show promise, they introduce trade-offs in complexity, cost, and security. Taken together, these findings suggest that IL is *unavoidable in principle* as a structural consequence of AMM rebalancing, yet *controllable in practice* in its realized magnitude through design and deployment choices. Such control typically reallocates costs to LP activity and rebalancing risk, trader frictions (fees/slippage), or explicit hedging premia. Accordingly, this section synthesizes common findings from the literature, evaluates mitigation strategies and their trade-offs, and outlines open questions.

A. Empirical and Theoretical Evidence

Across models and data, IL emerges as an *inherent* feature of CFMMs rather than a corner case [13], [20]. Canonical formulas for CPMMs and their multi-asset extensions quantify IL as the underperformance of LP portfolios relative to HODL. Moreover, the magnitude of IL is closely related to token volatility [13], [15]. Alternative benchmarks such as LVR sharpen this view by turning arbitrage-related losses into a measurable and forecastable cost [20].

Empirically, IL has been widely reported in Uniswap v2/v3, Balancer, and Curve. In some cases, fee income may offset IL, though not always [12], [15], [23]. When fees are reinvested into the invariant or adapt dynamically, small trades can generate impermanent gain (IG); under these conditions, pool returns align more closely with observed outcomes [16], [27].

The evidence points to two main conclusions: IL is ubiquitous in deployed AMMs, and fees can partially hedge IL but do not universally eliminate it. Beyond LP profitability, several works also emphasize that IL has ecosystem-level spillovers.

1) System-wide spillovers: slippage and leveraged LPing: Beyond LP returns, CFMM mechanics induce value transfers among traders, arbitrageurs, and liquidity providers. In constant-product settings, Labadie [42] links trader slippage to LP outcomes by showing that slippage can be interpreted as the mirror image of LP impermanent gain, reframing execution costs as a redistribution channel.

Composability can further amplify IL exposure through leverage. Wang et al. [43] study leveraged liquidity provision (e.g., borrowing to take larger LP positions, sometimes recursively), which increases fee income potential but also magnifies divergence-driven drawdowns and liquidation risk. Using Alpha Homora v2 position data, they show that leveraged borrowers are broadly exposed to IL through interactions with external AMM pools, and that performance depends strongly on timing and price moves. These findings reinforce that IL is not only an LP-level phenomenon, but can also act as a channel through which shocks propagate across interconnected DeFi protocols.

B. Mitigation Strategies and Their Trade-offs

The spillovers above motivate mitigation not only as a question of LP performance, but also as a design choice that redistributes costs and risks across market participants. We divide mitigation strategies into two categories and examine the challenges associated with each.

Table VII
SUMMARY OF TRADE-OFFS IN IL MITIGATION STRATEGIES

Aspect/Challenge	(a) Protocol Design Changes	(b) Financial-Engineering Overlays
Core Mechanism	Modifies an AMM's core logic, such as pricing curvature, liquidity shape, or fee policies.	Uses external financial instruments (e.g., options, variance swaps) to hedge an LP's exposure to IL.
Primary Advantage	Can reconfigure participant incentives and reduce IL at its source within the protocol.	Provides a formal, quantifiable hedge that converts IL exposure into a standard derivative payoff.
Primary Cost	Indirect costs, which may manifest as *higher slippage* for traders or *lower pool utilization* and fee income for LPs.	Direct costs, primarily through *paying option premiums* or fees for other derivative contracts.
Implementation Hurdle	Faces significant *governance and implementation hurdles*, as changes require protocol-level approval and development.	Can be rendered impractical or too costly by *trading frictions* like wide bid–ask spreads, slippage, and limited market depth.
Key Risks Introduced	Introduces potential *security, latency, and liveness risks*, especially in oracle-assisted designs. Often *shifts risk* rather than eliminating it.	Exposes LPs to *basis and model risk* (the hedge may not perfectly track IL) and *execution risk*.
External Dependencies	Relies on on-chain governance processes and, in some cases, the reliability and security of external *price oracles*.	Relies on the existence of liquid and efficient *off-chain or on-chain derivatives markets*.

(a) **Protocol design changes.** AMM variants adjust pricing curvature, liquidity shape, and fee policy to reduce IL or reallocate value [27], [34]–[36], [39]–[41]. *Trade-off:* Such modifications require protocol governance approval and careful evaluation of incentive effects. While some changes may reduce IL, they can increase slippage for traders or reduce utilization and fee income for LPs, effectively shifting costs rather than removing risk. Moreover, oracle-assisted designs [39], [40] introduce additional on-chain complexity and may entail latency, liveness, and security risks.

(b) **Financial-engineering overlays.** Static or semi-static option hedges, variance/weighted-variance swaps, and explicit IL protection claims convert LP exposure into standard derivative payoffs [28]–[31], [33]. *Trade-off:* Hedging strategies require paying option premia, are exposed to basis and model risk, and may rely on off-chain liquidity or periodic rebalancing. In thin options markets, trading frictions (wide bid–ask spreads, slippage, and limited depth) can reduce or even outweigh the protection implied by theory.

Both categories offer partial mitigation but come with non-trivial drawbacks. Protocol-level changes can reconfigure incentives yet face governance and implementation hurdles, while financial overlays provide formal hedges but introduce direct costs and execution risks. In practice, effective mitigation may require combining these approaches rather than relying on a single solution. Table VII summarizes the core concepts and trade-offs of these strategies.

C. Open Questions and Research Directions

Our analysis highlights several promising research directions:

Q1. Mitigation on existing deployments. What is the maximal IL reduction achievable on existing protocols? Can pool-specific controllers provably improve LP utility under realistic flow and gas constraints?

Q2. Affordable protection. Can on-chain IL insurance be priced and collateralized so that the expected cost is small relative to fees in normal regimes? What low-cost hedging strategies can limit extreme IL while allowing LPs to retain the majority of their fee income? Are weighted-variance proxies implementable with on-chain primitives?

Q3. Welfare and distributional effects. A key question is how the costs of reducing IL are redistributed among market participants. Do traders bear these costs through higher slippage, do arbitrageurs lose revenue through lower rents, or do LPs bear them through reduced utilization and fee income? Quantifying these distributional effects under different AMM designs can provide valuable guidance for protocol governance.

Q4. Customizable AMMs and modular mitigation. A promising direction lies in customizable AMMs, ex-

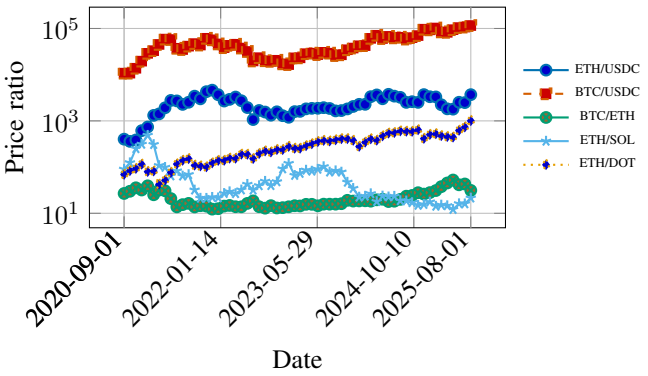

Figure 7. Price ratios of multiple coins between September 2020 and August 2025.

emplified by the proposed *Uniswap v4* [71]. The protocol introduces "hooks" [72], external smart contracts that execute at specific points in trading or liquidity provision actions. Hooks provide a framework for implementing pool-level IL mitigation strategies, such as adaptive fees that adjust to volatility or automated rebalancing of concentrated liquidity to keep positions in range. This architecture transforms mitigation from a protocol-wide design choice into a modular, pool-specific feature. An open question is what new vectors for economic exploitation or systemic risk emerge from this modularity, and whether formal methods can guarantee the safety and incentive-compatibility of arbitrary hooks.

Q5. Preventing IL indirectly. Impermanent loss can be increased by shocks outside the AMM, notably liquidations in lending protocols [45], [47]. For instance, in the Aave platform [44], a single liquidation moved ETH by -6.95% on Uniswap [46]. An open question is which non-direct events (liquidations, oracle failures, bridge de-pegs) affect IL the most, and whether they can be mitigated.

Q6. Risk modeling via currency correlation. Prior work [24] shows that volatility is a primary driver of IL, and pools with highly correlated or stable assets (e.g., stablecoin pairs) experience far less IL. This motivates modeling IL risk using historical price correlation between pool tokens. Figure 7 illustrates the diversity of these dynamics over five years. An open question is whether these observed correlations can be formalized into a predictive risk framework, which could inform adaptive fee mechanisms that apply higher fees when correlation-based risk is elevated.

Q7. Impact of Blockchain Infrastructure. How does blockchain infrastructure affect realized IL? In particular, block times, sequencing rules, and mempool transparency may shape how quickly arbitrage corrects price deviations and how much value is extracted during that process. Faster block times and higher throughput could shorten the arbitrage window and thereby reduce realized loss-versus-rebalancing (LVR) for LPs [21]. Likewise, private or protected mempools may reduce the MEV-related component of LP losses by limiting front-running and sandwich attacks during rebalancing. An open question is whether these infrastructure-level differences lead to systematically different IL outcomes across chains and execution environments.

VIII. Conclusion

This paper systematizes the literature on impermanent loss, the underperformance of liquidity provider portfolios compared to the buy-and-hold benchmark when asset prices diverge. The paper combines the theoretical impermanent loss models, empirical evidence, and mitigation proposals. We find that while approaches such as protocol redesigns and financial-derivative strategies can mitigate impermanent loss, each comes with trade-offs. Open questions remain regarding mitigation within existing deployments, the design of affordable protection instruments, and risk models that integrate volatility and correlation. Future research is essential to ensure that liquidity provision in DeFi remains both profitable for LPs and sustainable for AMMs.

Acknowledgments

AI content disclosure. OpenAI, Google Gemini, and Grammarly were used solely for language editing (grammar, clarity, and LaTeX/formatting suggestions). The authors are responsible for the technical content and for verifying the final manuscript.

References

[1] F. Schär, "Decentralized finance: On blockchain-and smart contract-based financial markets," *FRB of St. Louis Review*, 2021.
[2] D. A. Zetzsche, D. W. Arner, and R. P. Buckley, "Decentralized finance," *Journal of Financial Regulation*, vol. 6, no. 2, pp. 172–203, 2020.
[3] J. R. Jensen, V. von Wachter, and O. Ross, "An introduction to decentralized finance (DeFi)," *Complex Systems Informatics and Modeling Quarterly*, no. 26, pp. 46–54, 2021.
[4] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Decentralized Business Review*, 2008.
[5] A. Kotzer, D. Gandelman, and O. Rottenstreich, "SoK: Applications of sketches and rollups in blockchain networks," *IEEE Transactions on Network and Service Management*, vol. 21, no. 3, pp. 3194–3208, 2024.
[6] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, 2014.
[7] Forbes. (2025) Decentralized finance (DeFi) market cap. https://www.forbes.com/digital-assets/categories/decentralized-finance-defi-crypto/?sh=e9200737493d.
[8] CoinGecko. (2025) Decentralized finance (DeFi) market cap today. https://www.coingecko.com/en/categories/decentralized-finance-defi#:~:text=The%20Decentralized%20Finance%20(DeFi)%20market%20cap%20today,services%20such%20as%20lending%2C%20borrowing%2C%20and%20trading.
[9] A. Kotzer, T. Azoulay, Y. Abels, A. Yaish, and O. Rottenstreich, "SoK: DeFi lending and yield aggregation protocol taxonomy, empirical measurements, and security challenges," *IEEE Transactions on Network and Service Management*, 2026.
[10] J. Chiu, E. Ozdenoren, K. Yuan, and S. Zhang, "On the fragility of DeFi lending," Bank of Canada Staff Working Paper, Tech. Rep., 2023.
[11] K. Saengchote, "Decentralized lending and its users: Insights from compound," *Journal of International Financial Markets, Institutions and Money*, vol. 87, p. 101807, 2023.
[12] S. Loesch, N. Hindman, M. B. Richardson, and N. Welch, "Impermanent loss in Uniswap v3," *arXiv preprint arXiv:2111.09192*, 2021.
[13] N. Tiruviluamala, A. Port, and E. Lewis, "A general framework for impermanent loss in automated market makers," *arXiv preprint arXiv:2203.11352*, 2022.
[14] A. A. Aigner and G. Dhaliwal, "Uniswap: Impermanent loss and risk profile of a liquidity provider," *arXiv preprint arXiv:2106.14404*, 2021.
[15] N. Boueri, "G3M impermanent loss dynamics," *arXiv preprint arXiv:2108.06593*, 2021.
[16] V. Roman, V. Gorgadze, and A. Barger, "The impact of the exchange fees on impermanent loss of liquidity providers for conservative automated market makers," *The Journal of The British Blockchain Association*, vol. 8, no. 2, 2025.
[17] A. Evans, "Liquidity provider returns in geometric mean markets," *arXiv preprint arXiv:2006.08806*, 2021.
[18] J. Aoyagi, "Liquidity provision by automated market makers," *Available at SSRN 3674178*, 2020.
[19] G. Angeris, T. Chitra, and A. Evans, "When does the tail wag the dog? Curvature and market making," *arXiv preprint arXiv:2012.08040*, 2020.
[20] J. Milionis, C. C. Moallemi, T. Roughgarden, and A. L. Zhang, "Quantifying loss in automated market makers," in *ACM CCS Workshop on Decentralized Finance and Security*, 2022.
[21] ——, "Automated market making and loss-versus-rebalancing," *arXiv preprint arXiv:2208.06046*, 2022.

[22] S. N. Cohen, M. Sabaté-Vidales, D. Šiška, and Ł. Szpruch, "Inefficiency of CFMs: hedging perspective and agent-based simulations," in *International Conf. on Financial Cryptography and Data Security (FC)*, 2023.

[23] R. Tangri, P. Yatsyshin, E. A. Duijnstee, and D. Mandic, "Generalizing impermanent loss on decentralized exchanges with constant function market makers," *arXiv preprint arXiv:2301.06831*, 2023.

[24] L. Heimbach, Y. Wang, and R. Wattenhofer, "Behavior of liquidity providers in decentralized exchanges," *arXiv preprint arXiv:2105.13822*, 2021.

[25] M. Hafner and H. Dietl, "Impermanent loss conditions: An analysis of decentralized exchange platforms," *arXiv preprint arXiv:2401.07689*, 2024.

[26] L. Heimbach, E. Schertenleib, and R. Wattenhofer, "Risks and returns of Uniswap v3 liquidity providers," in *ACM Conference on Advances in Financial Technologies (AFT)*, 2022.

[27] I. Lebedeva, D. Umnov, Y. Yanovich, I. Melnikov, and G. Ovchinnikov, "Dynamic fee for reducing impermanent loss in decentralized exchanges," *arXiv preprint arXiv:2506.03001*, 2025.

[28] J. Deng, H. Zong, and Y. Wang, "Static replication of impermanent loss for concentrated liquidity provision in decentralised markets," *Operations Research Letters*, vol. 51, no. 3, pp. 206–211, 2023.

[29] A. Khakhar and X. Chen, "Delta hedging liquidity positions on automated market makers," *arXiv preprint arXiv:2208.03318*, 2022.

[30] A. M. González, J. I. Sequeira, and A. Dembling, "Pool value replication (CPM) and impermanent loss hedging," *arXiv preprint arXiv:2503.21967*, 2025.

[31] M. Fukasawa, B. Maire, and M. Wunsch, "Weighted variance swaps hedge against impermanent loss," *Quantitative Finance*, vol. 23, no. 6, pp. 901–911, 2023.

[32] M. Fukasawa, "Volatility derivatives and model-free implied leverage," *International Journal of Theoretical and Applied Finance*, vol. 17, no. 01, p. 1450002, 2014.

[33] A. Lipton, V. Lucic, and A. Sepp, "Unified approach for hedging impermanent loss of liquidity provision," *Digital Finance*, pp. 1–49, 2025.

[34] H. J. Kim, S. Choi, Y. T. Yoon, and S. Yoo, "Impermanent loss and gain of automated market maker smart contracts," *Authorea Preprints*, 2022.

[35] G. M. Lee and H. J. Kim, "Impermanent loss mitigation for decentralized exchanges through optimization," *International Journal of Industrial Engineering: Theory, Applications and Practice*, vol. 31, no. 6, 2024.

[36] T. Tran, D. A. Tran, and T. Nguyen, "Order book inspired automated market making," *IEEE Access*, vol. 12, pp. 36 743–36 763, 2024.

[37] V. Nadkarni, J. Hu, R. Rana, C. Jin, S. Kulkarni, and P. Viswanath, "Adaptive curves for optimally efficient market making," 2024. [Online]. Available: https://arxiv.org/abs/2406.13794

[38] L. Baggiani, M. Herdegen, and L. Sánchez-Betancourt, "Optimal dynamic fees in automated market makers," *arXiv preprint arXiv:2506.02869*, 2025.

[39] D. team, "Dodo," https://docs.dodoex.io/en/home/what-is-dodo, 2020.

[40] W. Chen, S. Chen, and P. Rozwood, "Improving capital efficiency and impermanent loss: Multi-token proactive market maker," *arXiv preprint arXiv:2309.00632*, 2023.

[41] N. Pattanaik, "Minimizing impermanent loss in asset-pair balancer pools," 2023.

[42] M. Labadie, "Impermanent loss and slippage in automated market makers (amms) with constant-product formula," *Available at SSRN 4053924*, 2022.

[43] Z. Wang, K. Qin, D. V. Minh, and A. Gervais, "Speculative multipliers on DeFi: Quantifying on-chain leverage risks," in *International Conference on Financial Cryptography and Data Security (FC)*, 2022.

[44] Aave Labs, "Aave protocol v3 documentation: Overview," https://aave.com/docs/developers/aave-v3/overview, 2022.

[45] ——, "Liquidations — Aave protocol documentation," https://aave.com/docs/developers/liquidations, 2022.

[46] K. Qin, J. Ernstberger, L. Zhou, P. Jovanovic, and A. Gervais, "Mitigating decentralized finance liquidations with reversible call options," in *International Conference on Financial Cryptography and Data Security (FC)*, 2023.

[47] T. Azoulay, U. Carl, and O. Rottenstreich, "Allowing blockchain loans with low collateral," in *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2023.

[48] E. Hertzog, G. Benartzi, and G. Benartzi, "Bancor protocol," *Continuous Liquidity for Cryptographic Tokens through their Smart Contracts. Available online: https://storage. googleapis. com/website-bancor/2018/04/01ba8253-bancor_protocol_whitepaper_en. pdf*, 2017.

[49] M. Bartoletti, J. H.-Y. Chiang, and A. Lluch-Lafuente, "A theory of automated market makers in DeFi," *Logical Methods in Computer Science*, vol. 18, no. 4, pp. 12:1–12:46, 2022.

[50] G. Angeris and T. Chitra, "Improved price oracles: Constant function market makers," in *ACM Conference on Advances in Financial Technologies (AFT)*, 2020.

[51] D. Pusceddu and M. Bartoletti, "Formalizing automated market makers in the lean 4 theorem prover," in *OpenAccess Series in Informatics (OASIcs)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024. [Online]. Available: https://arxiv.org/abs/2402.06064

[52] Y. Sokolik, M. Nassar, and O. Rottenstreich, "Age-aware fairness in blockchain transaction ordering for reducing tail latency," *IEEE Trans. Netw.*, vol. 33, no. 2, pp. 807–822, 2025.

[53] L. Heimbach and R. Wattenhofer, "Eliminating sandwich attacks with the help of game theory," in *ACM Asia Conference on Computer and Communications Security (ASIACCS)*, 2022.

[54] L. Zhou, K. Qin, C. F. Torres, D. V. Le, and A. Gervais, "High-frequency trading on decentralized on-chain exchanges," in *IEEE Symposium on Security and Privacy (SP)*, 2021.

[55] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *IEEE symposium on security and privacy (SP)*, 2020.

[56] K. Qin, L. Zhou, and A. Gervais, "Quantifying blockchain extractable value: How dark is the forest?" in *IEEE Symposium on Security and Privacy (SP)*, 2021.

[57] K. Gogol, C. Killer, M. Schlosser, T. Bocek, B. Stiller, and C. Tessone, "SoK: Decentralized finance (DeFi)–Fundamentals, taxonomy and risks," *arXiv preprint arXiv:2404.11281*, 2024.

[58] B. Llacer Trotti, W. Tang, R. El-Azouzi, G. Fanti, and D. S. Menasché, "Strategic analysis of just-in-time liquidity provision in concentrated liquidity market makers," in *ACM Conference on Advances in Financial Technologies (AFT)*, 2025.

[59] "Uniswap v2 core," https://uniswap.org/whitepaper.pdf, 2020.

[60] "Balancer protocol," https://balancer.fi/whitepaper.pdf, 2020.

[61] "Curve finance whitepaper," https://curve.fi/files/curve-stableswap-paper.pdf, 2020.

[62] "Uniswap v3 core," https://uniswap.org/whitepaper-v3.pdf, 2021.

[63] X. Wan and H. Adams, "Just-in-time liquidity on the Uniswap protocol," SSRN, 2022. [Online]. Available: https://ssrn.com/abstract=4382303

[64] A. Capponi, R. Jia, and B. Zhu, "The paradox of just-in-time liquidity in decentralized exchanges," *arXiv preprint arXiv:2311.18164*, 2023.

[65] L. R. Glosten and P. R. Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of financial economics*, vol. 14, no. 1, pp. 71–100, 1985.

[66] D. Simon, "Kalman filtering," *Embedded systems programming*, vol. 14, no. 6, pp. 72–79, 2001.

[67] F. Black and M. Scholes, "The pricing of options and corporate liabilities," *Journal of political economy*, vol. 81, no. 3, pp. 637–654, 1973.

[68] P. Glasserman, *Monte Carlo methods in financial engineering*. Springer, 2004, vol. 53.

[69] S. L. Heston, "A closed-form solution for options with stochastic volatility with applications to bond and currency options," *The review of financial studies*, vol. 6, no. 2, pp. 327–343, 1993.

[70] Deribit, "Deribit bitcoin options data," https://www.deribit.com/.

[71] H. Adams, S. Reynolds, M. Toda, M. Salem, A. Adams, A. Henshaw, D. Robinson, N. Zinsmeister, W. Pote, and E. Williams, "Uniswap v4 core," https://app.uniswap.org/whitepaper-v4.pdf, 2024.

[72] Uniswap Labs, "Hooks in Uniswap v4," 2025. [Online]. Available: https://docs.uniswap.org/contracts/v4/concepts/hooks