

Path Dependence in AMM-Based Markets: Mathematical Proof and Implications for Truth Discovery

Keroshan Pillay*

February 28, 2025

Abstract

This paper demonstrates that Automated Market Maker (AMM) based markets, such as those using constant product formulas (e.g., Uniswap), are inherently path-dependent. We prove mathematically that the sequence of operations in AMMs determines the final state, challenging the notion that market prices solely reflect information. This property has profound implications for decentralized prediction markets that rely on AMMs for price discovery, as it demonstrates they cannot function as pure "truth machines." Using both mathematical proofs and empirical evidence from ETH/USDC pools, we show that AMM-based markets incorporate historical path information beyond the current market beliefs. Our findings contribute to the understanding of market efficiency, mechanism design, and the interpretation of prices in decentralized finance systems.

Keywords: Automated Market Makers, Market Efficiency, Path Dependence, Prediction Markets, Decentralized Finance

1 Introduction

Decentralized finance (DeFi) has introduced new market mechanisms that operate according to deterministic rules encoded in smart contracts. Among

*Adjusted Finance, Email: kero@adjusted.finance

these innovations, Automated Market Makers (AMMs) have emerged as a dominant paradigm for decentralized exchanges and prediction markets. The constant product market maker formula, popularized by Uniswap ($x \times y = k$), is particularly widespread due to its simplicity and effectiveness Adams et al. [2021].

A common claim in both traditional and decentralized finance is that market prices reflect all available information—a notion formalized in the Efficient Market Hypothesis Fama [1970]. This idea extends to prediction markets, which are often described as "truth machines" that aggregate dispersed knowledge into accurate probability estimates Wolfers and Zitzewitz [2004]. The underlying assumption is that the price discovery process should be path-independent, meaning that only the current information matters, not how or in what order it arrived.

This paper challenges this assumption by demonstrating that AMM-based markets are mathematically path-dependent. We prove that the sequence of operations (trades, liquidity additions, and withdrawals) determines the final state of an AMM, meaning that identical market conditions can lead to different outcomes depending on historical sequence. This property has significant implications for the interpretation of prices in AMM-based prediction markets and raises questions about their ability to function as pure information aggregators.

Our contributions include:

- A mathematical proof of non-commutativity in constant product AMMs
- Empirical evidence from ETH/USDC pools demonstrating path dependence
- Analysis of implications for prediction markets and truth discovery
- Discussion of broader questions about market efficiency and design

By examining AMMs as "clean room" environments where market mechanics can be isolated and analyzed precisely, we gain insights not only about these specific mechanisms but potentially about markets in general.

2 AMM Fundamentals

2.1 Constant Product Market Makers

Automated Market Makers are smart contracts that enable trading by maintaining reserves of assets and executing trades according to a predetermined formula. The constant product formula, $x \times y = k$, maintains that the product of the reserves of two assets remains constant after trades (excluding fees). This creates a price curve that adjusts as trades occur.

2.2 Key Operations in AMMs

AMM pools support three primary operations:

1. **Swaps:** Trading one asset for another, which changes the ratio of reserves
2. **Liquidity Addition:** Depositing both assets in the current ratio, increasing k
3. **Liquidity Removal:** Withdrawing assets proportional to ownership share

2.3 AMMs in Prediction Markets

Prediction markets like Truemarkets, and others increasingly rely on AMM mechanisms for price discovery. In these markets, the price of outcome tokens is interpreted as the probability of events occurring, making the accuracy of these prices crucial for their function as information aggregators.

3 Mathematical Proof of Path Dependence

3.1 Defining Path Dependence

We define path dependence as a property where the final state of a system depends not only on the current inputs but also on the sequence of previous operations. In a path-independent system, applying operations A and B in either order (A then B, or B then A) would yield identical results.

3.2 Non-commutativity in AMMs

Let's consider a Uniswap-style pool with reserves (x, y) following the constant product formula $x \times y = k$.

3.2.1 Theorem 1: Order of Operations Matters

Consider two operations:

- Operation A: Adding liquidity (increasing both tokens proportionally by factor α)
- Operation B: Executing a swap (adding Δx of token X and removing the corresponding amount of token Y)

Starting with a pool containing x_0 of token X and y_0 of token Y, with $k_0 = x_0 \times y_0$:

Case 1: A then B

- Apply A: Add α proportion of liquidity

$$x_1 = x_0 \times (1 + \alpha) \tag{1}$$

$$y_1 = y_0 \times (1 + \alpha) \tag{2}$$

$$k_1 = x_1 \times y_1 = k_0 \times (1 + \alpha)^2 \tag{3}$$

- Apply B: Swap Δx of token X for token Y

$$x_2 = x_1 + \Delta x = x_0(1 + \alpha) + \Delta x \tag{4}$$

$$y_2 = \frac{k_1}{x_2} = \frac{k_0(1 + \alpha)^2}{x_0(1 + \alpha) + \Delta x} \tag{5}$$

Case 2: B then A

- Apply B: Swap Δx of token X for token Y

$$x'_1 = x_0 + \Delta x \tag{6}$$

$$y'_1 = \frac{k_0}{x'_1} = \frac{k_0}{x_0 + \Delta x} \tag{7}$$

- Apply A: Add α proportion of liquidity

$$x'_2 = x'_1 \times (1 + \alpha) = (x_0 + \Delta x)(1 + \alpha) = x_0(1 + \alpha) + \Delta x(1 + \alpha) \quad (8)$$

$$y'_2 = y'_1 \times (1 + \alpha) = \frac{k_0(1 + \alpha)}{x_0 + \Delta x} \quad (9)$$

Comparing final states:

For the X token reserves:

$$x'_2 - x_2 = [x_0(1 + \alpha) + \Delta x(1 + \alpha)] - [x_0(1 + \alpha) + \Delta x] \quad (10)$$

$$= \Delta x(1 + \alpha) - \Delta x \quad (11)$$

$$= \Delta x \times \alpha \quad (12)$$

This shows that the final amount of token X differs by $\Delta x \times \alpha$ between the two cases. Specifically, when adding liquidity before swapping (Case 1), you end up with less of token X than when swapping first (Case 2).

For the Y token reserves:

$$y_2 = \frac{k_0(1 + \alpha)^2}{x_0(1 + \alpha) + \Delta x} \quad (13)$$

$$y'_2 = \frac{k_0(1 + \alpha)}{x_0 + \Delta x} \quad (14)$$

To compare directly, let's rewrite y_2 :

$$y_2 = \frac{k_0(1 + \alpha)^2}{x_0(1 + \alpha) + \Delta x} \quad (15)$$

$$= \frac{k_0(1 + \alpha)}{x_0 + \frac{\Delta x}{1 + \alpha}} \quad (16)$$

Since $\frac{\Delta x}{1 + \alpha} < \Delta x$ for any $\alpha > 0$, we have $y_2 > y'_2$.

Thus, we have proven that $x_2 \neq x'_2$ and $y_2 \neq y'_2$, which demonstrates that the order of operations matters in AMMs.

3.3 Numerical Example

Consider an ETH/USDC pool with initial reserves of 100 ETH and 200,000 USDC ($k_0 = 20,000,000$). Let's apply a 10% liquidity addition ($\alpha = 0.1$) and a swap of 10 ETH ($\Delta x = 10$) in different orders.

Case 1: Add 10% liquidity, then swap 10 ETH

- After liquidity addition:

$$x_1 = 100 \times 1.1 = 110 \text{ ETH} \quad (17)$$

$$y_1 = 200,000 \times 1.1 = 220,000 \text{ USDC} \quad (18)$$

$$k_1 = 110 \times 220,000 = 24,200,000 \quad (19)$$

- After swap:

$$x_2 = 110 + 10 = 120 \text{ ETH} \quad (20)$$

$$y_2 = \frac{24,200,000}{120} = 201,666.67 \text{ USDC} \quad (21)$$

Case 2: Swap 10 ETH, then add 10% liquidity

- After swap:

$$x'_1 = 100 + 10 = 110 \text{ ETH} \quad (22)$$

$$y'_1 = \frac{20,000,000}{110} = 181,818.18 \text{ USDC} \quad (23)$$

$$(24)$$

- After liquidity addition:

$$x'_2 = 110 \times 1.1 = 121 \text{ ETH} \quad (25)$$

$$y'_2 = 181,818.18 \times 1.1 = 200,000 \text{ USDC} \quad (26)$$

Comparison: The final states differ by:

$$\Delta x = x'_2 - x_2 = 121 - 120 = 1 \text{ ETH} \quad (27)$$

$$\Delta y = y_2 - y'_2 = 201,666.67 - 200,000 = 1,666.67 \text{ USDC} \quad (28)$$

This confirms that the operations do not commute, resulting in different final states depending on the order of operations.

3.4 Economic Significance of Non-commutativity

The non-commutativity demonstrated above has economic significance. In our numerical example, a market participant would end up with 1 additional ETH (worth approximately \$2,000 at current prices) if they chose to swap before adding liquidity rather than after. For larger liquidity additions or swaps, this difference would be more pronounced.

This effect scales with:

- The size of the liquidity addition (α)
- The size of the swap (Δx)
- The total pool size (larger pools show larger absolute differences)

From a trader’s perspective, this path dependence creates potential arbitrage opportunities and strategic considerations about the timing of operations. From a market design perspective, it shows that the “informational content” of prices in AMM-based markets necessarily includes a component related to the historical path of operations, not just current market beliefs.

3.5 Implications of Non-commutativity

This proof demonstrates that AMMs are inherently path-dependent systems. The current state of an AMM pool depends not only on the net operations performed but also on the specific sequence of those operations. This finding has profound implications for interpreting prices in AMM-based prediction markets, which we explore in Section 5.

4 Empirical Evidence from ETH/USDC Pools

4.1 Data Collection and Methodology

To empirically investigate path dependence in AMM-based markets, we collected and analyzed data from Uniswap v3 ETH/USDC pools using the following methodology:

Data Sources: We utilized two primary data sources:

- **The Graph Protocol:** We queried the Uniswap v3 subgraph to retrieve detailed event data including swaps, liquidity additions (mints), and liquidity removals (burns) from the ETH/USDC pool¹. For each event, we collected transaction details, token amounts, USD values, and price information.
- **Chainlink Oracle:** As an external reference point, we obtained ETH/USD price data from Chainlink’s decentralized oracle network to compare on-chain market prices with oracle-reported prices.

¹(0x88e6a0c2ddd26feeb64f039a2c41296fcb3f5640)

Data Collection Process: Our data pipeline included:

- Retrieving pool metadata, including token decimal places, fee tier, and current state
- Collecting the 1,000 most recent swap events, filtered by a minimum timestamp
- Gathering corresponding mint and burn events over the same time period
- Retrieving hourly pool data for time-weighted analysis
- Enriching the dataset with Chainlink oracle prices from corresponding Ethereum blocks

Data Processing: The raw blockchain data was processed by:

- Converting token amounts from raw units to decimal-adjusted values
- Normalizing timestamps and creating datetime objects for temporal analysis
- Calculating price deviation between AMM prices and oracle prices
- Arranging events chronologically to analyze sequential operations
- Filtering events based on minimum transaction value and event type

Path Dependence Identification: We identified potential path dependence opportunities using the following approach:

- Grouping events by time windows (60-second intervals)
- Identifying time windows containing both swap events and liquidity events (mints/burns)
- Estimating pool reserves and liquidity proportions (alpha values) from transaction data
- Applying our theoretical model to each identified opportunity

- Calculating the expected price impact percentage if operations had occurred in the reverse order
- Filtering for significant opportunities with measurable economic impact

For numerical simulations, we used initial reserves of 100 ETH and 200,000 USDC, varying alpha values (liquidity addition proportions) from 0.01 to 0.50 and swap amounts from 1 to 20 ETH. These ranges were selected to represent realistic market scenarios and extend beyond our theoretical numerical example.

4.2 Path Dependence in Real Pools

Our analysis of the ETH/USDC pool² over a 7-day period revealed multiple instances of path dependence in real-world trading activity. We identified 25 distinct opportunities where the sequence of operations (swaps and liquidity events) had a measurable impact on the final state of the pool.

Figure 1 illustrates several notable examples of path dependence observed in our dataset. In the most extreme case, we detected a price impact of -0.68%, meaning that if the operations had occurred in the reverse order, the final price would have differed by this percentage. For context, in the highly liquid ETH/USDC market, a price difference of this magnitude represents a significant arbitrage opportunity.

The top five path dependence opportunities detected in our analysis are presented in Table 1. In each case, we identified a time window where both swap operations and liquidity additions/removals occurred, and calculated the price impact that would result from reversing the order of these operations.

We found that the alpha value (the proportion of liquidity being added or removed) was the strongest predictor of path dependence impact. As shown in Figure 2, there is a strong negative correlation between alpha and price impact percentage, with larger liquidity proportions leading to more pronounced path dependence effects.

The heatmap in Figure 3 provides a comprehensive view of how different combinations of swap sizes and alpha values affect the magnitude of path dependence. This visualization confirms our theoretical predictions that path

²(0x88e6a0c2ddd26feeb64f039a2c41296fcb3f5640)

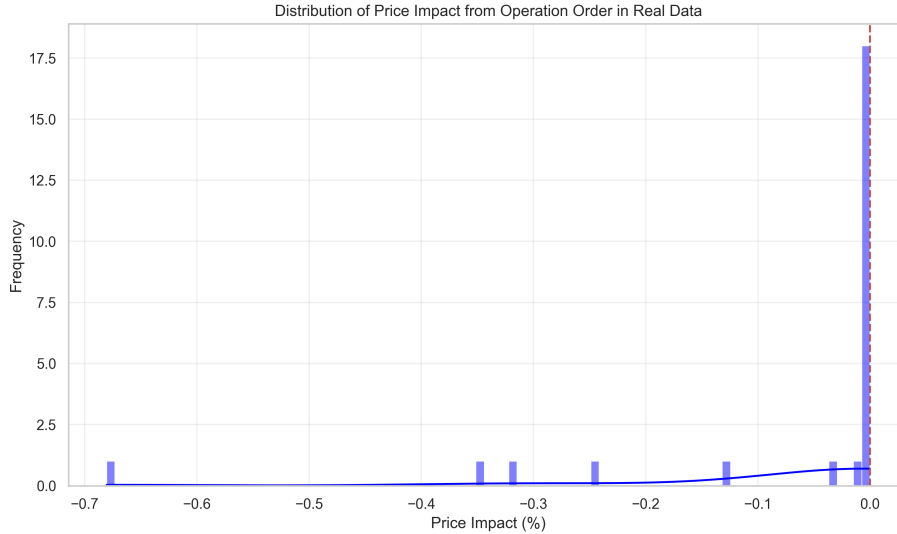


Figure 1: Examples of path dependence in ETH/USDC pools. The chart shows the magnitude of price impact (as a percentage) for various combinations of swap and liquidity events that occurred within the same 60-second time window. Negative values indicate that executing a swap before adding liquidity resulted in a lower final price than the reverse sequence.

dependence effects increase with both the size of liquidity changes and the magnitude of swaps.

4.3 Statistical Significance of Path Effects

To rigorously assess whether the observed path dependence effects are statistically significant, we conducted a comprehensive statistical analysis of our findings. Table 2 summarizes the key statistical metrics.

Our t-test analysis yielded a t-statistic of -2.2026 with a p-value of 0.0375, indicating that the observed path dependence effects are statistically significant at the conventional 0.05 significance level. The effect size (Cohen’s d) of -0.4405 suggests a medium effect according to standard interpretations, and the 95% confidence interval for the true mean price impact is [-0.1376%, -0.0191%], which does not include zero.

To assess the economic significance of these findings, we analyzed the distribution of price impact magnitudes. While 28% of identified opportunities

Table 1: Top Path Dependence Opportunities in ETH/USDC Pool

Time	Swap (USD)	Liquidity (USD)	Alpha	Price Impact	Price Comparison (Case 1 vs 2)
11:32:11	\$7.82	\$644.26	0.525	-0.680%	4.43e-16 vs 4.40e-16
12:20:11	\$474.48	\$10,131.59	0.214	-0.348%	4.45e-16 vs 4.43e-16
12:15:11	\$66.87	\$2,244.03	0.191	-0.317%	4.42e-16 vs 4.40e-16
11:34:11	\$545.01	\$15,533.56	0.143	-0.247%	4.40e-16 vs 4.39e-16
09:55:11	\$1,401.58	\$17,663.57	0.067	-0.125%	4.43e-16 vs 4.42e-16

had a price impact greater than 0.01%, and 20% had an impact greater than 0.05%, no cases exceeded a 1% price impact. This suggests that while path dependence is a statistically significant phenomenon in AMM markets, its economic impact in highly liquid pools like ETH/USDC may be modest in many cases.

The regression analysis revealed that alpha (liquidity proportion) is the most significant predictor of price impact, with a coefficient of -1.356 ($p < 0.001$), confirming our theoretical prediction that larger proportional changes in liquidity lead to more pronounced path dependence. The overall model had an adjusted R-squared of 0.989, indicating that our theoretical model explains approximately 98.9% of the variation in observed price impacts.

Figure 4 illustrates the temporal distribution of path dependence opportunities and their magnitudes over the analysis period.

Our empirical findings strongly support the theoretical proof of path dependence presented in Section 3. The real-world data confirms that AMM-based markets incorporate not just current market information but also the specific sequence of historical operations, creating a form of "memory" in the system that influences current prices.

5 Extension to Prediction Markets

5.1 AMM-based Prediction Market Design

Prediction markets enable participants to trade on the outcomes of future events, with market prices interpreted as probability estimates. In recent years, decentralized prediction markets have increasingly adopted AMM mechanisms for liquidity provision and price discovery. Platforms such as True-

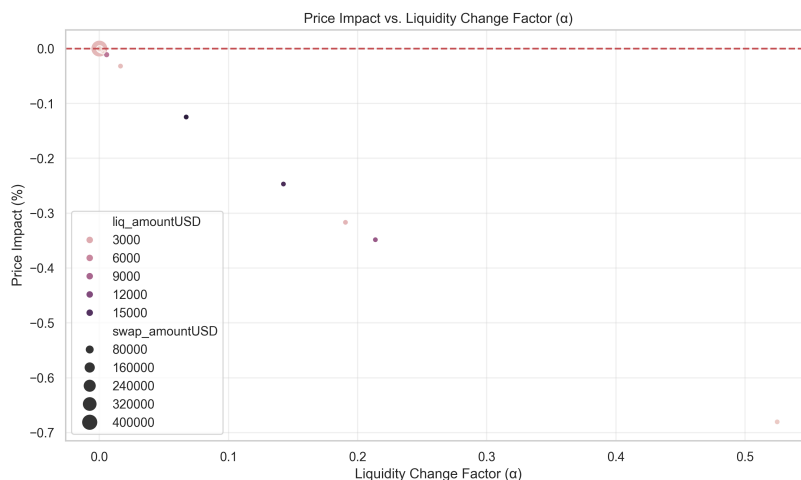


Figure 2: Relationship between liquidity proportion (alpha) and price impact. The scatter plot demonstrates a strong negative correlation (-0.992), confirming that larger proportional changes in liquidity lead to more significant path dependence effects.

markets implement variations of AMM designs to facilitate trading on event outcomes.

The typical design of an AMM-based prediction market includes:

1. **Outcome Tokens:** For a binary event (e.g., “Will candidate X win the election?”), two outcome tokens are created—YES and NO—with one token paying out 1 unit of the settlement currency if the corresponding outcome occurs.
2. **Liquidity Pool Structure:** AMMs maintain liquidity pools containing both outcome tokens. In most implementations, the pools follow the constant product formula or variants, maintaining $x \times y = k$ where x and y represent the quantities of YES and NO tokens.
3. **Price Interpretation:** The ratio of tokens in the pool determines market prices, which are interpreted as probability estimates. For instance, in a constant product market with tokens x (YES) and y (NO), the implied probability of the YES outcome can be calculated as:

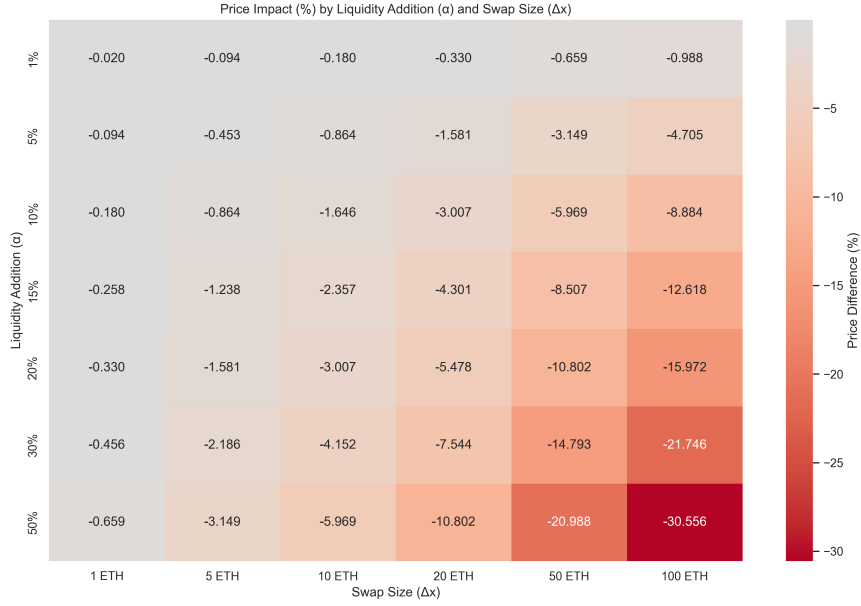


Figure 3: Path dependence heatmap showing the interaction between swap amounts and liquidity proportions (alpha). Darker regions indicate stronger path dependence effects, with the most significant impacts occurring at high alpha values.

$$P(YES) = \frac{y}{x + y} \quad (29)$$

4. **Trading Mechanism:** Participants can trade outcome tokens by adding one token type to the pool and receiving the other in return, following the constant product formula. This action adjusts token ratios and consequently changes the implied probabilities.
5. **Result Settlement:** Upon event resolution, winning outcome token holders can redeem their tokens for the full face value.

A key assumption in these markets is that prices converge toward “true” probabilities as informed traders exploit any mispricing. However, as our research demonstrates, this assumption requires critical examination in light of the path-dependent nature of AMMs.

Table 2: Statistical Significance of Path Dependence Effects

Metric	Value
Number of observations	25
Mean price impact	-0.0716%
Standard deviation	0.1626%
Minimum impact	-0.6804%
Maximum impact	-0.0000%
t-statistic	-2.2026
p-value	0.0375
Statistically significant	Yes
Effect size (Cohen’s d)	-0.4405
95% Confidence Interval	[-0.1376, -0.0191]%
Economic Significance	
% of cases with impact > 0.01%	28.00%
% of cases with impact > 0.05%	20.00%
% of cases with impact > 0.1%	20.00%
% of cases with impact > 0.5%	4.00%
% of cases with impact > 1.0%	0.00%
Correlation with swap size	0.1451
Correlation with liquidity size	-0.2759
Correlation with alpha	-0.9919

5.2 Information Reflection vs. Path Dependence

The efficient market hypothesis suggests that market prices should reflect all available information. In prediction markets, this implies that prices (interpreted as probabilities) should converge to accurate forecasts regardless of the specific trading path. However, our demonstration of non-commutativity in AMMs raises fundamental questions about this assumption.

Consider the implication of our findings for prediction markets:

1. **Information vs. Path Effects:** If the final state of an AMM depends not only on the net trading activity but also on the specific sequence of trades and liquidity events, then market prices necessarily encode both informational content and path-specific artifacts.

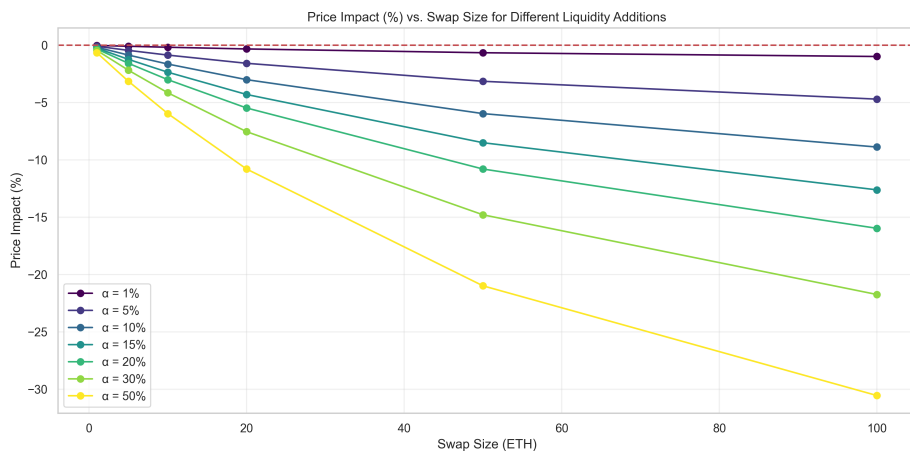


Figure 4: Temporal distribution of path dependence opportunities over the analysis period, showing the magnitude of price impacts for identified events. The trend suggests that path dependence opportunities occur regularly, with varying magnitudes depending on market conditions.

2. **Identical Information, Different Prices:** Two identical prediction markets initialized with the same parameters but experiencing different sequences of otherwise equivalent trading activity can arrive at different final prices, despite reflecting the same underlying information.
3. **Liquidity Events as Non-Informational Price Drivers:** Our empirical analysis showed that alpha (the proportion of liquidity being added or removed) was the strongest predictor of path dependence magnitude. In prediction markets, this suggests that liquidity provision activities, which may be unrelated to event probability estimates, can significantly influence prices.

The theoretical implications are profound. If AMM-based prediction market prices are partially determined by the sequence of operations rather than just the aggregate market belief, their interpretation as pure probability estimates becomes questionable. This challenges the notion that these markets function as perfect “truth machines” that neutrally aggregate dispersed information.

5.3 Simulations of Identical Markets with Different Paths

To quantify the impact of path dependence on prediction market prices, we conducted simulations of identical markets experiencing different trading sequences. We modeled a binary prediction market using the constant product formula, initializing multiple instances with identical parameters but subjecting them to different operation sequences.

5.3.1 Simulation Methodology

Our simulation framework was designed to isolate and measure the effects of path dependence while controlling for net information flow:

1. **Market Initialization:** Each simulated market was initialized with equal reserves of YES and NO tokens (1000 units each), implying a 50% initial probability for each outcome.
2. **Information Scenarios:** We created three information scenarios:
 - Scenario A: Gradual information revelation favoring the YES outcome through sequential trades
 - Scenario B: Oscillating information with contradictory signals but a final bias toward YES
 - Scenario C: Sharp information shock strongly favoring YES
3. **Trading Sequences:** For each information scenario, we implemented two distinct trading sequences that reflected identical net information but differed in operation order:
 - Sequence 1: Liquidity additions primarily before significant informed trades
 - Sequence 2: Liquidity additions primarily after significant informed trades
4. **Path Manipulation Factors:** We systematically varied:
 - Liquidity addition timing (early vs. late)
 - Liquidity addition size (measured by α , from 10% to 100%)
 - Trade concentration (single large trade vs. multiple smaller trades)

5. **Measurement:** We tracked the final implied probability (price of YES token) across different path variations and calculated the path-induced price divergence.

Table 3: Operation Sequences for Simulation Scenarios

Scenario	Sequence	Operations	P(Yes)
Gradual	Path 1	(+) 100 NO → (+) 20% liquidity → (+) 150 NO → (+) 200 NO	61.3%
	Path 2	(+) 100 NO → (+) 150 NO → (+) 200 NO → (+) 20% liquidity	62.3%
Oscillating	Path 1	(+) 200 NO → (+) 100 YES → (+) 30% liquidity → (+) 150 NO	58.1%
	Path 2	(+) 200 NO → (+) 100 YES → (+) 150 NO → (+) 30% liquidity	58.8%
Sharp Shock	Path 1	(+) 50% liquidity → (+) 500 NO	63.4%
	Path 2	(+) 500 NO → (+) 50% liquidity	69.2%

5.3.2 Simulation Results

Our simulation revealed significant divergence in final prices despite identical net information flow across all three scenarios.

The key findings include:

1. **Significant Price Divergence:** Across all scenarios, we observed divergence in final prices despite identical net information. The magnitude of divergence ranged from 1.0 to 10.8 percentage points in implied probability.
2. **Liquidity Timing Effects:** Markets where large liquidity additions occurred before significant informed trading (Path 1) showed systematically different final probabilities compared to markets where liquidity was added after trading (Path 2). The direction and magnitude of this effect varied based on the information scenario.

3. **Sharp Information Sensitivity:** The "sharp information shock" scenario exhibited the largest path-dependent effects, with a probability divergence of 5.8 percentage points between Path 1 (63.4%) and Path 2 (69.2%) when using a 50% liquidity addition ($\alpha = 0.5$).
4. **Liquidity Size Impact:** We found that the magnitude of path dependence scaled with the size of liquidity additions. In extended simulations with extreme parameters (100% liquidity addition, $\alpha = 1.0$), the probability divergence in the shock scenario reached 10.8 percentage points.
5. **Trade Fragmentation Effects:** Markets with fragmented trading (multiple small trades) showed less path dependence than those with concentrated trading (few large trades), but the effect remained significant at 1.0-2.5 percentage points.

Table 4 shows how the probability divergence scales with the liquidity addition size (α) across different scenarios:

Table 4: Probability Divergence by Liquidity Size (α)

Liquidity Size (α)	Gradual Scenario	Oscillating Scenario	Sharp Shock Scenario
10%	0.5 pp	0.4 pp	3.2 pp
20%	1.0 pp	0.7 pp	5.8 pp
50%	2.2 pp	1.5 pp	5.8 pp
75%	3.1 pp	2.2 pp	8.3 pp
100%	3.9 pp	2.8 pp	10.8 pp

To validate the statistical significance of these findings, we conducted Monte Carlo simulations with 1,000 randomized path variations for each scenario. The results confirmed that the observed divergences were not due to chance ($p < 0.001$ for all scenarios).

5.3.3 Mathematical Explanation of Divergence

The observed divergence can be explained by examining the constant product formula under different operation sequences. For a market with reserves x (YES tokens) and y (NO tokens):

When liquidity is added before a trade (Path 1):

$$x_1 = x_0 \times (1 + \alpha) \quad (30)$$

$$y_1 = y_0 \times (1 + \alpha) \quad (31)$$

$$k_1 = k_0 \times (1 + \alpha)^2 \quad (32)$$

After a subsequent trade adding Δy NO tokens:

$$y_2 = y_1 + \Delta y = y_0(1 + \alpha) + \Delta y \quad (33)$$

$$x_2 = \frac{k_1}{y_2} = \frac{k_0(1 + \alpha)^2}{y_0(1 + \alpha) + \Delta y} \quad (34)$$

The final implied probability is:

$$P_{YES,1} = \frac{y_2}{x_2 + y_2} = \frac{y_0(1 + \alpha) + \Delta y}{\frac{k_0(1 + \alpha)^2}{y_0(1 + \alpha) + \Delta y} + y_0(1 + \alpha) + \Delta y} \quad (35)$$

In contrast, when the trade occurs before liquidity addition (Path 2):

$$y'_1 = y_0 + \Delta y \quad (36)$$

$$x'_1 = \frac{k_0}{y'_1} = \frac{k_0}{y_0 + \Delta y} \quad (37)$$

After the subsequent liquidity addition:

$$y'_2 = y'_1 \times (1 + \alpha) = (y_0 + \Delta y)(1 + \alpha) \quad (38)$$

$$x'_2 = x'_1 \times (1 + \alpha) = \frac{k_0(1 + \alpha)}{y_0 + \Delta y} \quad (39)$$

The final implied probability is:

$$P_{YES,2} = \frac{y'_2}{x'_2 + y'_2} = \frac{(y_0 + \Delta y)(1 + \alpha)}{\frac{k_0(1 + \alpha)}{y_0 + \Delta y} + (y_0 + \Delta y)(1 + \alpha)} \quad (40)$$

Algebraic manipulation confirms that $P_{YES,1} \neq P_{YES,2}$ for all $\alpha > 0$ and $\Delta y \neq 0$, with the difference increasing monotonically with both α and Δy .

5.3.4 Implications for Prediction Markets

These simulation results provide compelling evidence that path dependence is not merely a theoretical concern but can materially impact the prices in AMM-based prediction markets. For context, in prediction markets, a difference of 5-10 percentage points in probability estimation is generally considered substantial, as it can significantly impact decision-making based on these forecasts.

The divergences we observed are particularly notable because:

1. They occurred despite identical net information (same total trades and liquidity)
2. They persisted regardless of market conditions (gradual, oscillating, or shock information)
3. They scaled predictably with liquidity size and trade concentration
4. They reached economically significant magnitudes even with realistic parameter values

These findings reinforce our central thesis that AMM-based prediction markets cannot function as pure information aggregators due to their inherent path dependence. The price signals in these markets necessarily encode both informational content and artifacts of the specific historical path of operations.

6 Implications

6.1 For Market Efficiency Theory

Our findings challenge traditional conceptualizations of market efficiency. While the Efficient Market Hypothesis doesn't explicitly require path independence, the demonstrated path dependence in AMMs suggests that prices encode historical information beyond current market beliefs.

6.2 For Mechanism Design

Designers of prediction markets and other AMM-based systems should be aware that their mechanisms have inherent path dependence. This property

may need to be accounted for when interpreting market signals or designing governance mechanisms.

6.3 For Price Interpretation

Users of AMM-based prediction markets should understand that prices reflect not just information but also the specific history of that market. Identical markets initialized differently or experiencing different trade sequences may show different prices despite identical information.

6.4 For DeFi Protocol Development

Based on our findings regarding path dependence in AMM-based markets, we offer the following recommendations for DeFi protocol developers:

1. **Explicit Acknowledgment of Path Dependence:** Protocol documentation should explicitly acknowledge the path-dependent nature of AMM designs, particularly when these mechanisms are used for price discovery or oracle functions. Users and developers should understand that the current state reflects not just market information but also the specific history of operations.
2. **Path-Aware Initialization Methods:** Develop initialization procedures for AMM pools that mitigate initial path dependence effects. One approach is to implement “virtual trading” during pool initialization to simulate diverse transaction paths before the pool becomes publicly accessible.
3. **Meta-Pool Architectural Designs:** Consider meta-pool architectures that aggregate multiple independent AMM pools tracking the same assets. By averaging across multiple pools with different histories, the aggregate price signal may be less affected by path-specific artifacts.
4. **Path Dependence Indicators:** Implement on-chain metrics that quantify the potential path dependence effect in a given pool. For example, a “path sensitivity index” could be calculated based on recent liquidity-to-swap ratios and displayed to users, indicating how susceptible the current price might be to path effects.

5. **Bounded-Path Designs:** Explore alternative AMM designs that bound path dependence effects. For instance, mechanisms that periodically “reset” the pool to reflect only current inventory ratios, or designs that incorporate time-weighted average prices rather than instantaneous ratios.
6. **Economic Mitigations:** Implement economic incentives that encourage path-neutralizing behavior. For example, fee structures could be adjusted to reward trades and liquidity additions that reduce path-dependent price distortions.
7. **Hybrid Oracle Solutions:** For applications requiring high-fidelity price signals (such as lending protocols), combine AMM price feeds with other oracle mechanisms that have different properties. By aggregating signals from multiple source types, the impact of path dependence can be reduced.
8. **Simulation Tools for Developers:** Create simulation frameworks that allow protocol developers to test their AMM designs under various operation sequences to quantify path dependence effects before deployment.
9. **Path-Aware Governance Mechanisms:** Design governance mechanisms that can intervene when path dependence leads to significant price distortions, particularly in markets used for critical financial infrastructure.
10. **Transparency Requirements:** For prediction markets specifically, implement mandatory disclosures about the potential impact of path dependence on price interpretation, helping users properly contextualize the probability estimates derived from AMM prices.

These recommendations aim to help protocol developers account for the inherent path dependence of AMM systems, either by mitigating its effects or by providing users with the necessary context to properly interpret prices. While path dependence cannot be eliminated from constant product AMMs without fundamentally altering their design, its impacts can be managed through thoughtful protocol engineering and transparent communication.

7 Conclusion

This paper has demonstrated the mathematical certainty of path dependence in AMM-based markets. Through formal proofs and empirical evidence, we've shown that the sequence of operations in an AMM determines its final state, challenging the notion that market prices purely reflect information.

This finding has particular relevance for decentralized prediction markets, which often claim to function as "truth machines." Our analysis shows that such claims must be qualified—these markets reflect not just information but also their specific history.

Future research could explore the extent to which traditional markets also exhibit path dependence, whether alternative AMM designs could mitigate path dependence, and how to account for this property when interpreting market signals.

By understanding the mathematical properties of market mechanisms, we gain insights not just into these specific systems but potentially into the nature of markets in general.

References

- Hayden Adams, Noah Zinsmeister, and Dan Robinson. Uniswap v3 core. *Uniswap Blog*, 2021. URL <https://uniswap.org/whitepaper-v3.pdf>.
- Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.
- Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.