CHAOS LABS

GMX V2 Genesis Risk Framework and Methodology

A rigorous analysis of GMX V2, focusing on market risk, protocol mechanism design, and risk parameter analysis on a per market basis.

GMX

July 25, 2023

GMX V2 Genesis Risk Framework and Methodology

Omer Goldberg
omer@chaoslabs.xyz

Shai Kritz
shai@chaoslabs.xyz

Yonatan Haimovich
haimo@chaoslabs.xyz

July 2023

Abstract

This study comprehensively analyzes the Version 2 (V2) configuration of the GMX platform, an Ethereum-based decentralized exchange, examining its implications for traders, liquidity providers, and the broader crypto market. This research endeavors to understand the new trading environment using Version 1 (V1) position data and simulate potential trading outcomes under various market conditions.

Using a historical approach, the paper investigates how price impact parameters affect traditional trading activities, quantifying the associated costs and prospective shifts in user behavior. Furthermore, it provides recommendations for the initial V2 market risk parameter configurations, seeking a balance between risk management and user engagement.

We then discuss the function and computation of funding fees, exploring a dual-slope model to mitigate imbalances and optimize capital utilization. The paper also scrutinizes the interplay between borrowing and funding fees, highlighting the intricate relationship underpinning open interest balancing mechanisms.

Chaos Labs' research utilizes Agent-Based Simulation (ABS), a renowned tool in domains like aeronautical engineering and defense, to conduct stress tests on GMX's smart contracts. Despite ABS's complexity and challenges in achieving congruence with experimental results, judicious design, tuning, and infrastructure architecture can mitigate these issues. The implementation of ABS, along with continuous data-driven adjustments, enables an accurate and risk-reflective model of GMX's trading environment.

Lastly, our study acknowledges the inherent constraints of the adopted methodology and proposes continuous iterations and adjustments based on future empirical evidence. This paper is essential for those seeking to understand the complexities of decentralized financial exchanges and the future trajectory of GMX's trading environment.

Contents

Disclaimer

This document is purely informational and does not constitute an invitation to acquire any security, an appeal for any purchase or sale, or an endorsement of any financial instrument. Neither is it an assertion of the provision of investment consultancy or other services by Chaos Labs Inc. References to specific securities should not be perceived as recommendations for any transaction, including buying, selling, or retaining any such securities. Nothing herein should be regarded as a solicitation or offer to negotiate any security, future, option, or other financial instrument or to extend any investment advice or service to any entity in any jurisdiction. The contents of this document should not be interpreted as offering investment advice or presenting any opinion on the viability of any security, and any advice to purchase, dispose of or maintain any security in this report should not be acted upon. The information contained in this document should not form the basis for making investment decisions.

While preparing the information presented in this report, we have not considered individual investors' specific investment requirements, objectives, and financial situations. This information does not account for the specific investment goals, financial status, and individual requirements of the recipient of this information, and the investments discussed may not be suitable for all investors. Any views presented in this report by us were prepared based on the information available when these views were written. Additional or modified information could cause these views to change. All information is subject to possible rectification. Information may rapidly become unreliable for various reasons, including market or economic changes.

Chapter 1

Overview

GMX V2 Prelaunch Risk Research

Chaos Labs has diligently worked alongside GMX to methodically research, evaluate, and refine the impending deployment of the GMX V2 application. Leveraging sophisticated Agent-Based Simulation Models, Chaos Labs pursues a delicate equilibrium between capital efficiency and risk exposure.

Nonetheless, given the early-stage status of V2, our immediate focus and priority are heavily tilted toward the security and safety of the protocol. We are committed to ensuring a smooth and safe phased launch, continuously monitoring trader activities and system performance, and laying a robust foundation upon which the protocol can reliably expand.

The ensuing set of genesis parameter recommendations is the fruit of our rigorous research process, which involved executing millions of Monte Carlo simulations. These simulations included a wide array of variables, such as diverse price trajectories, liquidity models, levels of network congestion, and volatility levels, among others, thereby testing the system's robustness in a myriad of potential scenarios.

The forthcoming sections provide a detailed analysis, a review of the results obtained from our simulations, and, ultimately, a careful examination of our recommendations. Each proposed solution will be analyzed, comprehensively discussing its merits and potential drawbacks. This rigorous process will ensure that the proposed solutions are academically sound, practically viable, and well-suited to meet the unique challenges and requirements of the GMX V2 protocol.

Please note that the preliminary parameter recommendations are formulated on the basis of a combination of economic-security-driven heuristics, an analysis of trader and liquidity provider (LP) behavior observed in V1, and in-depth research on empirical trading patterns found in centralized exchanges (CEX), alongside other empirical evidence. Post the launch of V2, Chaos Labs will closely monitor the platform usage and iteratively refine our models to favor analysis of agent behavior within the V2 environment increasingly. Consequently, these models and parameter recommendations will undergo further refinement and optimization, informed by these real-world observations.

Chapter 2

GMX V2 Economic Analysis

1 Dynamic Fee Structures and Comprehensive Pricing Evaluation

GMX V2 introduces a novel architecture explicitly designed to mitigate the risk exposures of liquidity providers by implementing segregated pools. These isolated pools serve to fine-tune control and management of exposure. Additionally, V2 accommodates synthetic pools, obviating the need to hold volatile assets for liquidity providers and traders. While GMX V1 delivered an enhanced trading user experience, primarily by obviating price impact on trades, it inadvertently heightened the liquidity providers' exposure to market skew, a phenomenon observed empirically.

In this vein, GMX V2 puts forth a unique fee structure. This design seeks to incentivize market neutrality and fortify defenses against market manipulation while concurrently ensuring competitive pricing for traders without resorting to explicitly restrictive open interest limits. Our analysis aims to evaluate the proposed structure's efficacy and establish the optimal parameter values to minimize liquidity providers' risk whilst optimizing the legitimate trading UX.

1.1 Synthetic Perpetual Exchanges

A Synthetic Perpetual Product is a type of financial instrument often found within the realm of crypto (decentralized and centralized finance), specifically designed to emulate the price of an underlying asset or index without necessitating the ownership of the underlying asset itself.

These products are termed "synthetic" because they comprise one or more derivatives that simulate an underlying asset. They're "perpetual" because, unlike futures contracts, they do not have an expiry date and can be held indefinitely.

The essence of a Synthetic Perpetual Product lies in its ability to mimic traditional financial instruments, such as futures, and facilitate unlimited exposure to various assets (like cryptocurrencies, commodities, stocks, indices, etc.) without the need for an actual asset exchange. Instead, these products are collateralized by a different asset, typically a stable cryptocurrency.

Synthetic Perpetual Products are traded on what's known as a Perpetual Swap Contract, a particular type of futures contract where, instead of settling on a predetermined date, the contract remains open until the trader decides to close it. These contracts often use a mechanism known as the "funding rate" to keep the synthetic asset's price pegged to the real asset's price.

The ability of Synthetic Perpetual Products to offer exposure to any asset, combined with the flexibility of no expiry, makes them an attractive financial instrument in the crypto and DeFi space. This is particularly useful in jurisdictions where certain assets may be difficult to access due to various restrictions.

1.2 What are Price Impact Fees?

Synthetic Perpetual Exchange products leverage the concept of price impact to emulate the dynamics seen in traditional order book-based exchanges. This emulation is essential due to the inherent difference in the operating mechanism between synthetic perpetual exchanges and traditional ones.

In traditional order book-based exchanges, market participants place buy or sell orders at specified prices, creating a "depth of market" where the bid-ask spread represents the difference between the highest price that a buyer is willing to pay (bid price) and the lowest price at which a seller is willing to sell (ask price). The dynamics of these exchanges involve price movement due to matching these orders, a process that influences the supply-demand balance and, ultimately, the asset price.

On the other hand, Synthetic Perpetual Exchange products use a mechanism that directly interacts with a liquidity pool rather than matching individual orders. In these systems, trades result in a direct price impact, moving the asset price depending on the trade size. Therefore, the price impact mechanism is employed to simulate the behavior of price changes as observed in an order-book model.

The concept of price impact allows Synthetic Perpetual Exchanges to reflect the immediate price change based on the trade's size. Larger trades significantly impact the price, emulating how larger orders would move the price in an order book model. This is akin to 'slippage' in a traditional order book context, where large market orders can end up executing at less favorable prices due to the depth of the order book.

The benchmark for the order-book spread in traditional exchanges is often set by market conditions, precisely the balance of supply and demand and the overall liquidity in the market. Lower liquidity or higher volatility usually results in wider spreads as market participants demand a larger premium for taking on additional risk. This spread functions as a transaction cost and can serve as a measure of market efficiency - tighter spreads typically indicate a more efficient market.

In Synthetic Perpetual Exchange products, the spread can be considered a component of the overall transaction cost, but the specifics will depend on the individual protocol's design and mechanisms.

2 The Implications of Price Impact Fees

The primary objectives of integrating price impact fees are multifaceted:

2.1 Constructing safeguards to deter price manipulation

In Synthetic Perpetual Markets, price impact fees serve as a crucial lever in deterring price manipulation exploits. These fees are transaction costs that traders incur when the size of their order impacts the market price of an asset. The fees increase with larger order sizes, effectively reducing the potential profit from the trade and making it cost-prohibitive to manipulate the market price by placing large orders.

The rationale behind the indispensability of these fees can be articulated by delineating their impacts across three distinct pillars. It's noteworthy that while these pillars may intersect in terms of their effects, they vary in the pathways they employ to achieve their respective objectives. Thus, each pillar encapsulates a unique perspective on how these fees reinforce market stability and fairness.

1. **Increasing Capital Requirements and Risk of Ruin for Potential Market Manipulators:** Price impact fees can deter attempts to manipulate market prices. In the absence of these fees, a trader with significant resources could place large buy or sell orders to artificially inflate or deflate prices and then profit from these price movements. Price impact fees make such tactics financially unviable, thus reducing the likelihood of price manipulation.

2. **Maintaining Market Stability:** Price impact fees contribute to market stability by discouraging sudden, large trades. Large trades can cause significant price volatility, leading to a less predictable and stable market. These fees thus help maintain a more orderly and predictable market environment.

3. **Preventing Flash Crashes and Price Spikes:** Price impact fees can help prevent flash crashes and sudden price spikes. Large sell orders can rapidly drive down prices (causing a flash crash), while large buy orders can cause a quick spike. Both situations can lead to cascading liquidations or margin calls. Price impact fees can prevent such scenarios by making large trades more expensive.

4. **Aiding Effective Price Discovery:** Price impact fees support the fair price discovery process. They discourage 'quote stuffing' or other strategies intended to skew the price discovery process, helping ensure prices accurately reflect the balance between supply and demand.

Therefore, introducing price impact fees in Synthetic Perpetual Markets acts as an effective risk management tool, preserving market integrity and preventing potential price manipulation exploits.

2.2 Maintaining a balanced Open Interest (OI)

Balanced open interest in financial derivatives like futures, options, or synthetic perpetual products is crucial to ensure the stability and health of the market.

Open interest generally refers to the total number of active or open contracts in the market, regardless of whether they are bought or sold. This figure does not count contracts that have already been settled or closed.

Maintaining balanced, open interest is critical for the health of the protocol. Below we list several critical reasons:

1. **Incentivizing Market Stability:** An unbalanced open interest can lead to price distortions. For instance, if most traders long on a particular derivative contract, a sudden price drop could trigger a cascade of liquidations, pushing the price down further. This could potentially lead to market instability.

2. **Liquidity:** Balanced, open interest generally equates to good market liquidity. If there's a balance between buyers and sellers, it's usually easier to open and close positions without significantly impacting the price. On the other hand, an imbalance in open interest can lead to liquidity issues, making it harder to enter or exit positions without causing substantial price swings.

3. **Risk Management:** For financial platforms offering these derivatives, a balanced, open interest helps manage risk. When open interest is heavily tilted to one side (either long or short), the platform could be exposed to significant risk if the market moves sharply in the opposite direction.

4. **Fair Price Discovery:** Balanced, open interest contributes to efficient and fair price discovery. When there's an equilibrium between buying and selling pressure, prices better reflect the true value of the underlying asset based on current market information.

   Balancing open interest is essential to managing derivative products to ensure efficient functioning, adequate liquidity, robust risk management, and fair price discovery.

2.3 Facilitating Balanced Pools by Charging Price Impact on Spot Swaps

V2 segregated markets of GMX present versatility, extending their utility beyond mere speculative instruments to function as AMMs for spot swaps. The imposition of price impact fees is integral to this enhanced functionality. Without these fees, markets might be exploited by arbitrageurs seeking liquid sources that offer minimal or nonexistent price impact. Consequently, these price impact fees play a pivotal role in augmenting the market functionality within GMX V2. Furthermore, they act as an instrumental lever in incentivizing the equilibrium of liquidity pools, ultimately contributing to market stability and efficiency.

2.4 Facilitating Delta-Neutral Markets via Price Impact Fees

The Price Impact function in GMX V2 is structured to incentivize market neutrality. It encourages traders to adopt positions that counterbalance the prevailing market skew. Traders opening positions that increase market skew are funding trading opening positions on the minority side. When a market is predominantly long, traders who opt for short positions receive a premium funded by those opening long positions. In contrast, if a market is skewed short, traders opting for long positions benefit from a discount, while those opening short positions incur a premium. This incentive structure nudges traders towards positions that restore market equilibrium, enhancing stability and decreasing the risk profile for liquidity providers.

This mechanism effectively fosters a high-frequency rebalancing incentive while imposing soft constraints on the pool's maximum exposure. This is accomplished by levying "Negative Price Impact" fees on trades that exacerbate skew expansion and disbursing "Positive Price Impact" fees on trades contributing to skew compression.

2.5 Emulating Orderbook Dynamics

GMX V2 employs a price impact function to emulate order book depth dynamics. It is important to note that while similar in its intended effect, V2 does not identically adopt order book mechanics. Instead, V2 applies fee amounts derived from relative and absolute market skews. This aims to achieve a similar effect to spread in Centralized Exchanges, derived from real-time orderbook dynamics. In contrast to CEX spreads, V2 fees are solely applied when a majority side position is opened. Opening minority side positions can introduce a premium or discount on the execution price, contingent on the trade direction (long or short) and the current market skew (long or short). We will denote a shorthand notation for variables as follows for reader clarification. We will name $initialOpenInterestImbalance\ as\ initOIImbalance$, $nextOpenInterestImbalance\ as\ nextOIImbalance, priceImpactExponent\ as\ piExponent$, $and\ priceImpactFactor\ as\ piFactor$

With our shorthand notations, we present the calculation of price impact with the subsequent formula:

$$Price\ Impact = initOIImbalance^{piExponent} * \frac{piFactor}{2} - nextOIImbalance^{piExponent} * \frac{piFactor}{2} \tag{2.1}$$

We utilize imbalance as a metric to measure the difference in worth between long and short tokens for swaps. To illustrate, consider the following scenario:

- Consider a pool with 10 long tokens, each valued at \$5000, and 50,000 short tokens, each worth \$1. This results in an equally balanced pool, with long and short tokens totaling \$50,000.
- $piExponent = 2$
- $piFactor = \frac{\$0.01}{50{,}000}$
- If a user deposits 10 long tokens, the pool now contains \$100,000 in long tokens and \$50,000 in short tokens. The change in imbalance is from \$0 to \$50,000.
- A negative price impact is then charged on the user's deposit. Calculated as:

$$\$0^2 \times \frac{\$0.01}{50,000} - \$50,000^2 \times \frac{\$0.01}{50,000} = -\$500$$

Regarding position actions (increase/decrease), the imbalance is calculated as the difference between long and short open interest.

The purpose of this price impact calculation is threefold:

1. Incentivizing a balanced distribution of tokens within pools,
2. Encouraging a balance of long and short positions,
3. Minimizing the risk of price manipulation.

Note that this computation depicts the user's trade price impact rather than the pool's. For instance, while a user's trade may have a 0.25% price impact, a subsequent minuscule trade may have a 0.5% price impact.

Price manipulation risks may still persist if positive and negative price impact values are similar. Hence, during periods of volatility or irregular price movements, the positive price impact should be set low.

For price impact on position increases/decreases, instead of deducting collateral when a negative price impact occurs, the position's entry/exit price is adjusted based on the price impact. For example:

- Suppose the oracle price of the index token is $5000 and a user opens a long position of size $50,000 with a negative price impact of 0.1%.
- The user's position size in USD is $50,000, and the size in tokens is

$$\frac{50,000}{5000} \times (100 - 0.1)\% = 9.99\ tokens$$

- This results in an entry price of

$$\frac{50,000}{9.99} = \ \$5005$$

for the position.

- The negative price impact is recorded as a quantity of index tokens in the pool.
- The user's pending PnL at this point would be

$$50,000 - \frac{9.99}{5000} = \$50$$

- Tokens in the position impact pool should be factored into pool value computation to offset this pending PnL.
- The overall impact on the pool is neutral, +$50 from the pending negative PnL due to price impact and -$50 from the 0.01 index tokens in the position impact pool worth $50.
- If the user closes the position with a negative price impact of 0.2%, the position impact pool would rise to 0.03 index tokens.
- The user receives (original position collateral - $150).
- The pool possesses an extra $150 of collateral, maintaining a net zero impact on the pool value due to the 0.03 index tokens in the position impact pool.

Price impact exponents and price impact factors are configured per market and can differ for spot and position actions.

3 Risk Methodology

3.1 Classification of Markets

GMX V2 introduces segregated markets, allowing optimal risk levers per market adjustments and tailoring parameters to each index asset. While future enhancements will fine-tune every pool individually and implement on-chain logic that dynamically updates parameters based on market conditions and usage, we have initially identified three distinct market categories for genesis parameters:

We employ different methodologies for each category, considering their unique market properties and associated risks.

Market Classification		
Blue Chip	Medium Market Cap	Medium Market Cap Synthetics
BTC, ETH	LINK, UNI, AVAX, ARB, SOL	DOGE, LTC

3.2 Blue Chip

Considering the expansive liquidity of BTC and ETH, the risk of price manipulation is relatively low. We aim to configure price impact parameters such that within an "acceptable" skew, the price impact fees are competitive with CeFi venues. An acceptable skew could, for example, be a 60%/40% skew in a certain direction. When skew exceeds this point, the market becomes less "liquid" - heavily penalizing trades in the direction of skew and incentivizing trades in the opposite direction. This dynamic is similar to a low liquidity scenario on a CeFi venue's perpetual market when there is a large deviation from the spot price.

Since the price impact function (2.1) depends on the nominal value rather than the relative skew, we need to estimate the total open interest size and calculate the maximum accepted imbalance. This task may pose a challenge as the pool size might be highly volatile at launch. To capture the range of possible states, we aim for a 55%-45% skew, bounded by an open interest cap of $64M USD on the majority side. Although this cap might appear restrictive, it serves as a safety measure for the market launch and can be adjusted as the market finds balance. If the cap is fully utilized, the acceptable imbalance would be approximately $12M USD. If the pool is less utilized than expected, for instance, $20M USD on the majority side, the acceptable imbalance is approximately $3.6M USD.

Our initial goal is to ensure attractive trading when the pool is in perfect balance. In this situation, the function (1) simplifies to:

$$Price\ Impact = -tradeSize^{piExponent} * \frac{piFactor}{2} \tag{2.2}$$

The curve below is constructed to meet the following requirements for a balanced pool:

Trades under $1M result in less than 1 bps price impact fees. Trades in the $1M-$3M range yield price impact fees that increase linearly from 0.9 bps/$1M to between 1 and 1.9 bps. Trades in the $3M-$10M range yield price impact fees that increase linearly from 1bps/$1M to between 1.9 and 8 bps. Trades larger than $10M yield price impact fees that increase linearly at a rate of 1.5bps/$1M. These requirements are designed to make medium to large (defined as larger than $10M USD) trades competitive with CEXs, while heavily penalizing larger trades. They can be adjusted according to the acceptable trade size. We then use linear regression to fit the (2) curve to the desired fees, as the illustration below shows.

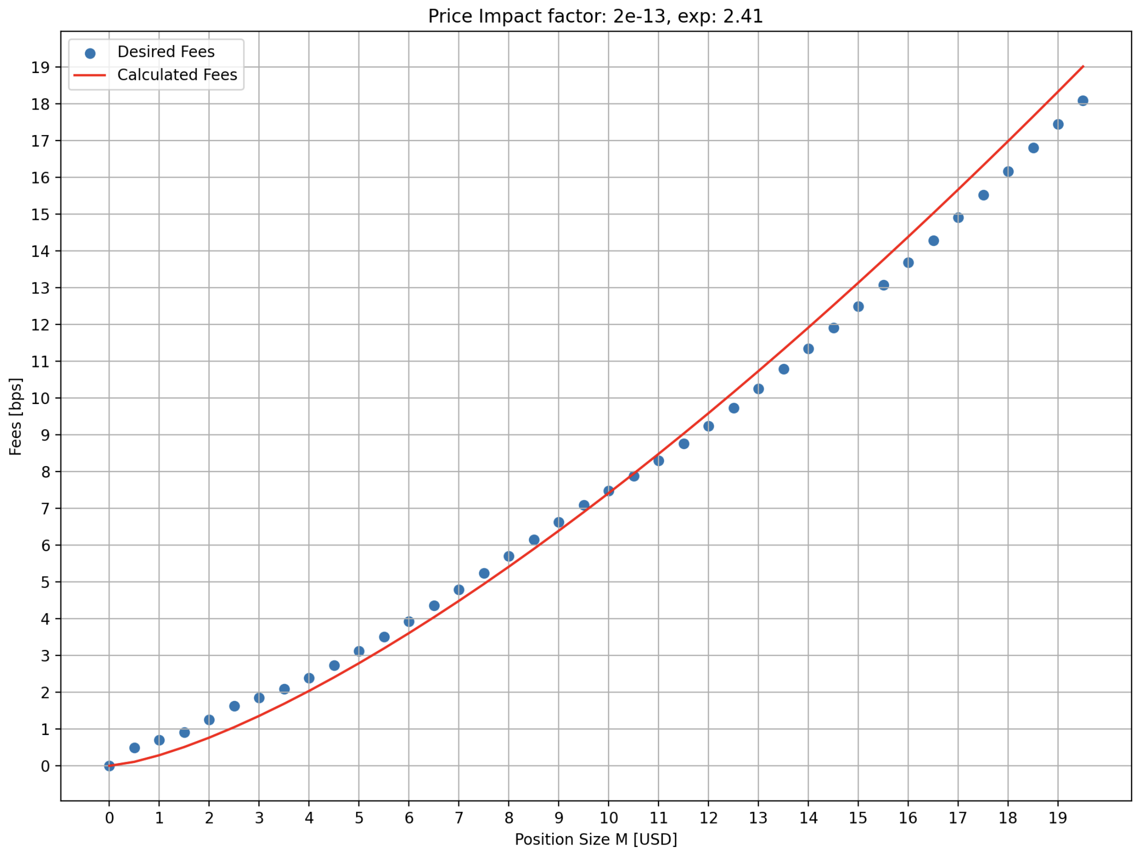

Figure 2.1: The blue scattered points are the desired fees chosen according to the requirements. The red plot best fits the GMX price impact model to the desired fees in MSE terms.

As observed above, the linear regression yields the following results:

$$\begin{aligned} piFactor &= 2 * 10^{-13} \\ piExponent &= 2.41 \end{aligned} \tag{2.3}$$

Since the precise target imbalance is unknown, we examine the fees imposed on trades in different states of initial imbalance:

price_impact_bps_factor_2e-13_exp_2.41

Position Size \ Initial Pool Imbalance	0.0 M	0.5 M	1.0 M	2.0 M	3.0 M	4.0 M	5.0 M	10.0 M	15.0 M
0.01 M	0.0	0.27	0.7	1.85	3.28	4.92	6.73	17.88	31.66
0.1 M	0.01	0.3	0.75	1.91	3.35	5.0	6.82	17.99	31.79
0.5 M	0.11	0.47	0.96	2.18	3.66	5.35	7.2	18.5	32.39
1.0 M	0.29	0.71	1.24	2.54	4.07	5.8	7.7	19.14	33.15
2.0 M	0.77	1.29	1.89	3.31	4.94	6.75	8.72	20.45	34.67
3.0 M	1.36	1.95	2.62	4.14	5.86	7.75	9.78	21.79	36.23
5.0 M	2.79	3.5	4.27	5.97	7.84	9.87	12.04	24.57	39.41
10.0 M	7.41	8.33	9.3	11.35	13.54	15.86	18.3	31.99	47.76
15.0 M	13.13	14.21	15.32	17.65	20.1	22.67	25.34	40.03	56.66
20.0 M	19.7	20.9	22.14	24.71	27.39	30.16	33.03	48.63	66.04

Figure 2.2: The table depicts price impact fees, in basis points units, as a function of position sizes and initial pool imbalances.

Our observations suggest that when the imbalance remains comparatively low (less than \$5M USD), the fees align reasonably well with the intended figures. However, when the imbalance reaches \$10M, even relatively small trades (less than \$0.5M USD) might be subject to substantial fees. Considering the permissible imbalance at full cap utilization is \$12M, we are inclined to propose more favorable pricing under such significant skew conditions. Therefore, we examine an alternative configuration:

$$\begin{aligned} piFactor &= 10^{-10} \\ piExponent &= 2 \end{aligned} \tag{2.4}$$

Setting the exponent to 2 yields a linear increase in relative fees when the pool is balanced while allowing a “soft” imbalance range where trading is competitive.

3.3 Medium Market Cap Assets

Within the domain of cryptocurrencies possessing a medium market capitalization, typically ranging between \$1B and \$10B USD, these digital assets generally demonstrate substantial liquidity and trading volume on Centralized Exchanges (CEXs). Nevertheless, these assets remain vulnerable to external factors capable of triggering severe market volatility. For instance, adverse news events, such as the publication of documents by the U.S. Securities and Exchange Commission (SEC) in June 2023, indicating the potential classification of specific assets as securities, can incite drastic price declines or flash crashes. Such occurrences jeopardize the robustness and stability of these digital assets, underlining the intricacies of cryptocurrency market dynamics. Consequently, this can create an environment conducive to price manipulation, rendering perpetual exchanges vulnerable to potential exploits. Specifically, low-margin requirements and guaranteed limit order executions could incentivize attackers to siphon off funds.

To counteract this threat, GMX V2 imposes price impact fees on potential exploiters to reduce the profitability of such attacks. The strategy is to calibrate these parameters such that the liquidity of GMX V2 never surpasses that of the external market. Towards this end, we employ simulations to assess the potential profitability of an attack. We model various scenarios using randomly generated price trajectories and emulate an attacker's behavior through strategically placed limit orders, to maximize profits through market manipulation. We elaborate further on the simulation environment, execution, and agent catalogs deployed in the Simulation Section.

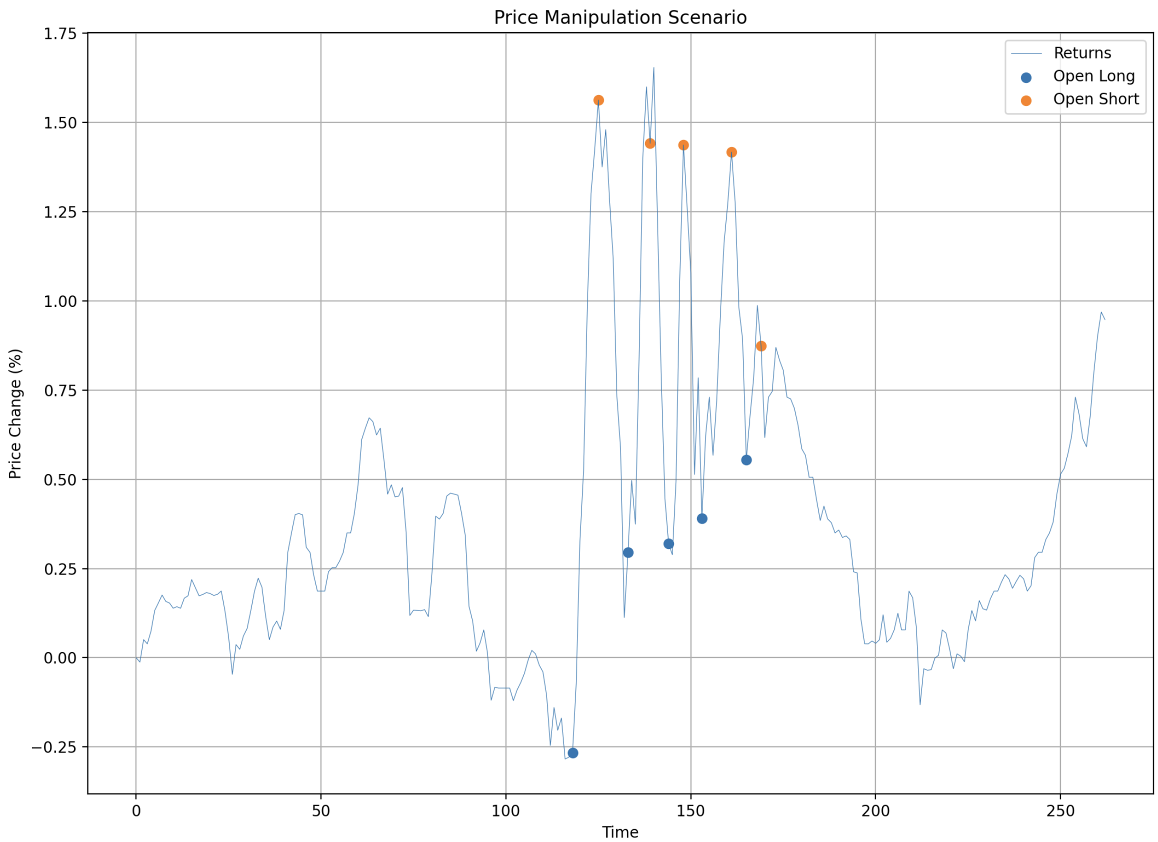

Figure 2.3: A time series of theoretical price manipulation.

Initial simulations reveal that when price impact parameters are zero, the attacker's Profit and Loss (PnL) stands at approximately $600K. To gain a comprehensive understanding, we expand our analysis to cover a range of parameter values to examine their effect on an attacker's PnL within GMX-Synthetics. It's important to note that this value does not account for the potential cost to the attacker of manipulating prices on CEXs.

To ascertain the optimal configuration, we estimate an attack's profitability and subsequently calculate the associated cost from the resultant PnL analysis. An attack is profitable if the attacker can secure a profit-to-cost ratio of at least 2:1, considering the maximum capital requirement. This criterion guides our selection of configurations to balance efficiency and security. The overarching goal is to deter potential attackers by minimizing their potential gains.

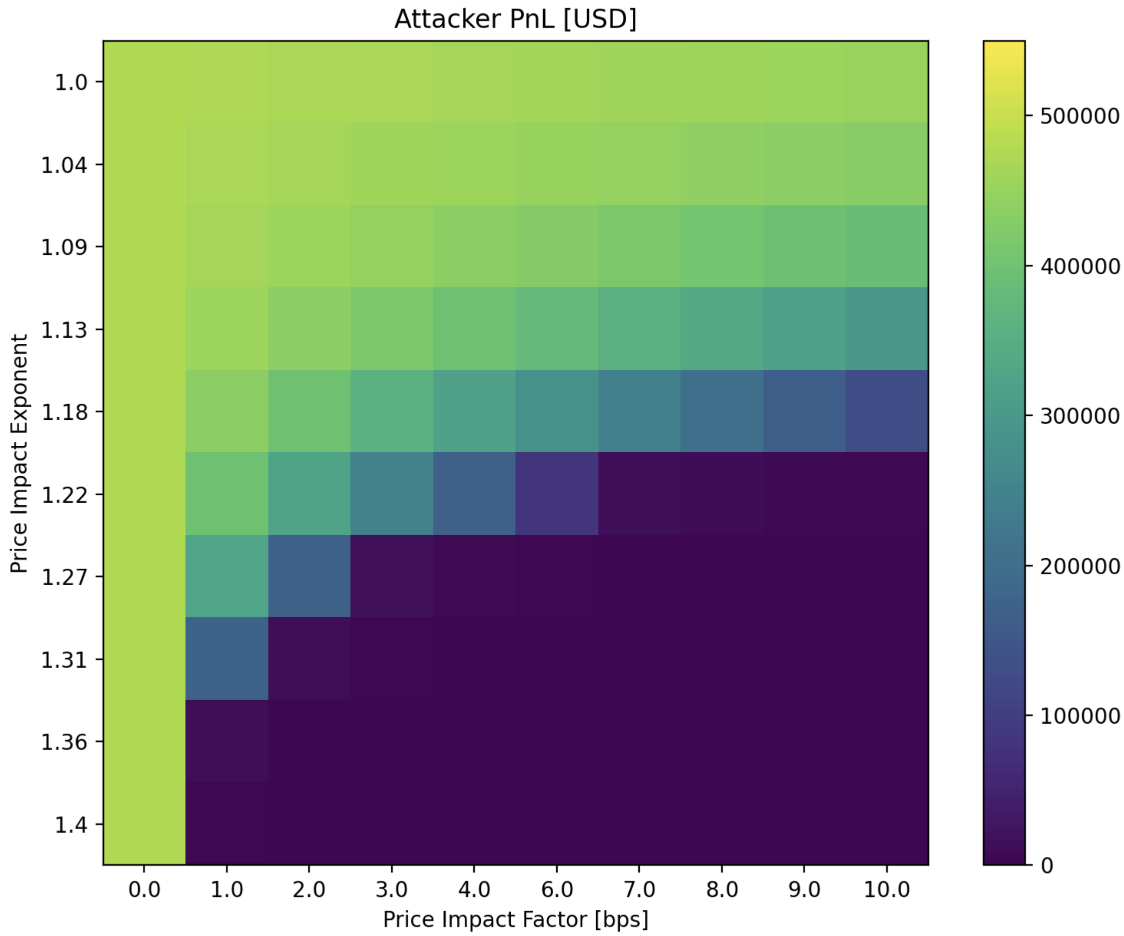

Figure 2.4: Quantifying the Attacker's Profit and Loss Across a Spectrum of Price Impact Parameters.

3.4 Assessing Attack Costs

Throughout our research process, we have devoted substantial resources to developing a sophisticated model referred to as PNLModel. This model has been engineered to predict with precision the potential yields that could emerge from a market manipulation attack. However, in light of this article's open and transparent disposition, we consciously refrain from sharing the nuanced details of the model. The unintentional misuse of such sensitive information could conceivably serve as a blueprint for orchestrating market manipulation attacks against various protocols.

In our pursuit of empirical validation, we subjected our theoretical model to rigorous backtesting against real-world market manipulations, such as the AVAX manipulation attack, thereby corroborating the model's accuracy. As we advance, we shall employ this particular exploit as a reference point for benchmarking.

The application of our PNL model establishes a lower bound for the cost of an attack at approximately $91K USD. To curb the profitability of an attack, we pinpoint a location on the heat map such that the profit is less than twice the cost of the attack. Formally, we define the following constraint:

$$PnL < 2 * AttackCost = \$182K\ USD \tag{2.5}$$

Kindly bear in mind that the inequality coefficient of 2 is amenable to modifications as required, thereby providing us with a lever to adjust the stringency or leniency of market risk configuration settings depending on the prevailing market conditions and risk appetite.

Running a parameter sweep simulation with our model yields the following Attacker PnL heatmap:

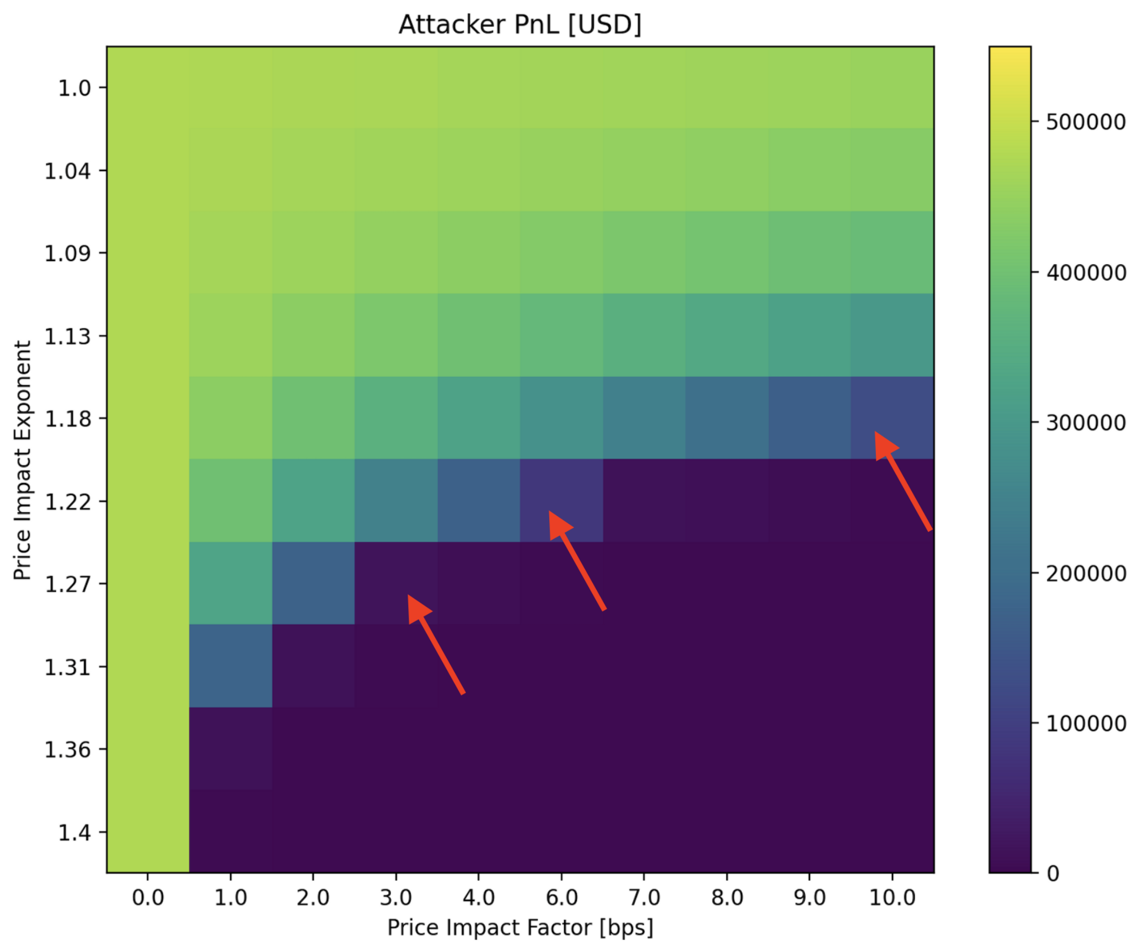

Figure 2.5: The presented graphic is a heatmap plot that displays the Attacker's Profit and Loss (PnL) as a function of the PriceImpactExponent and PriceImpactFactor. The grid cells highlighted in red represent coordinates that satisfy our previously defined constraint concerning the attacker's PnL.

To broaden our analysis to encompass other assets, we adopt an extrapolation technique based on the ratio of the trading volume and order book depth during the attack to the trading volume immediately preceding the attack. This comparative ratio serves as an instrumental variable in our extrapolation procedure, enabling us to predict the potential implications of analogous attacks on different assets. This methodology yields a benchmarking instrument that aids in comprehending and addressing potential vulnerabilities across diverse asset landscapes.

3.5 Establishing Positive Price Impact Parameters

The implementation of Positive Price Impact serves as a potent incentive for pool rebalancing. However, in potential attack scenarios, it could inadvertently enhance the profitability for the attacker, as they could receive compensation when closing positions and reverting the pool to its original state. Therefore, it is imperative that the negative price impact exceeds the positive price impact, thereby ensuring that an attacker's Profit and Loss (PnL) is insufficient to cover incurred costs. To this end, we initially establish:

$$negativePriceImpactFactor = 2 * positivePriceImpactFactor$$

Subsequent to this initial setting, we re-conduct the analysis. It's worth noting that this initial factor setting may be overly conservative, potentially levying higher fees than necessary given the inherent incentive for arbitrageurs to counterbalance the positive price impact before an attacker can close their position. Consequently, it is crucial that we gather empirical data on such arbitrage opportunities in a production environment before relying substantially on this mechanism.

3.6 Deciphering Interleaved Orders

An ill-intentioned actor could endeavor to circumvent the impact of price impact fees through a strategy of interleaving orders. This tactic comprises alternating between initiating long and short positions in increments. For instance:

1. Open a $1k long position for a 0.02% price impact.
2. Open a $1k short position for a 0.02% price impact.
3. Open a $1k long position for a 0.02% price impact.
4. Open a $1k short position for a 0.02% price impact

Implementing this strategy enables users to build long and short positions while bearing minimal price impact, as the pool's imbalance remains minimal throughout the process. Assuming a user successfully constructs a $10 million long position and a $10 million short position in this manner, they manage to manipulate the price upwards by 5%. If the short position only holds 2% of the position size as collateral, then the short position would be liquidated at the cost of:

$$\$10 million \times 2\% = \$200,000$$

Simultaneously, the long position accrues a profit of:

$$\$10 million \times 5\% = \$500,000$$

This long position could be closed with zero or a positive price impact post-liquidation since it would contribute to the re-equilibrium of the pool's balance.

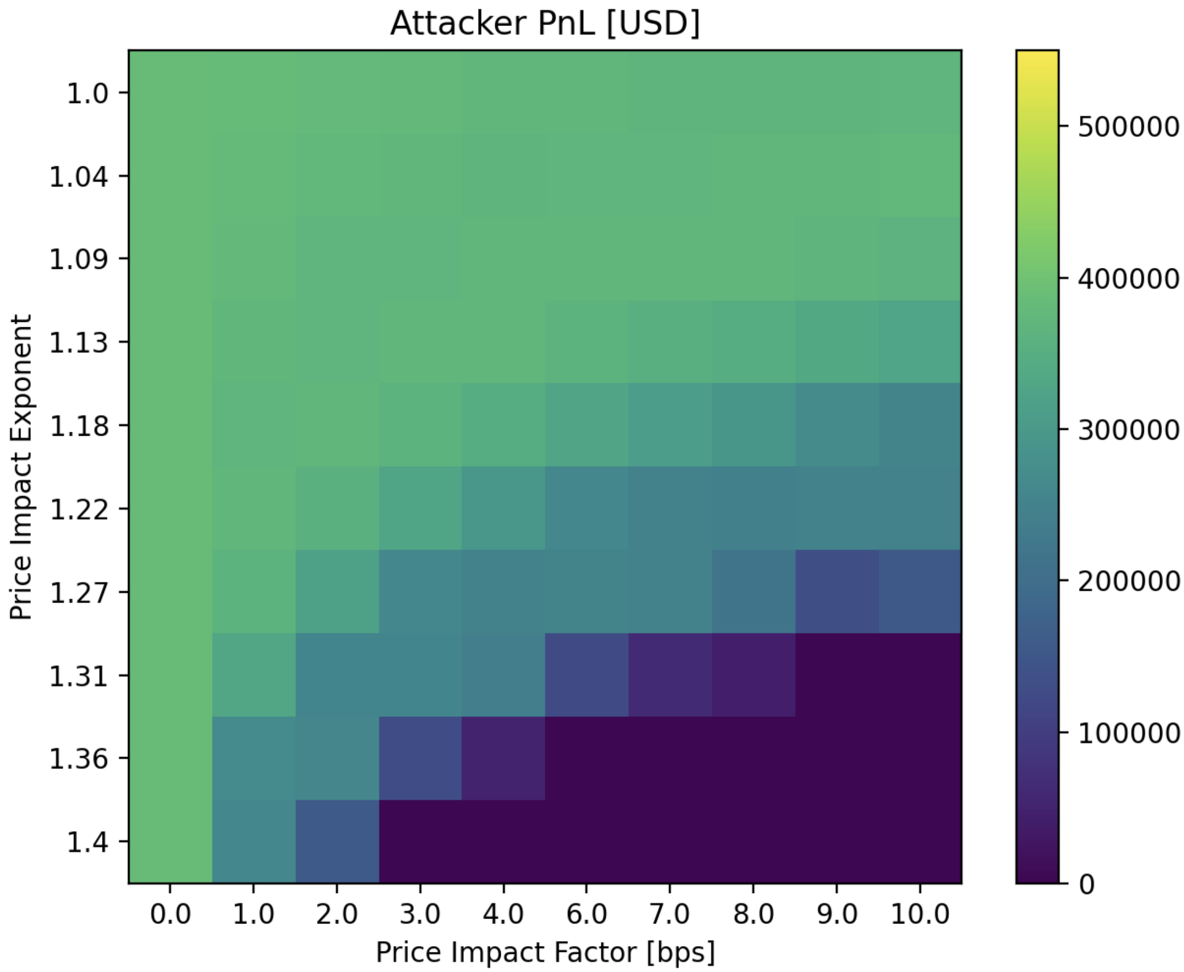

Figure 2.6: Quantifying the Attacker's Profit and Loss Across a Spectrum of Price Impact Parameters in the Context of Interleaved Orders.

3.7 Mitigating Interleaved Attacks with minCollateralFactorForOpenInterest-Multiplier

To curb this exploitation, the $minCollateralFactorForOpenInterestMultiplier$ is designed to be set such that the maximum allowed leverage decreases as the total open interest of the pool increases. This ensures that high-leverage positions cannot be utilized to manipulate the system and decrease the price impact amount. The following heatmap exhibits the attacker PnL when integrating $minCollateralFactorForOpenInterestMultiplier$.

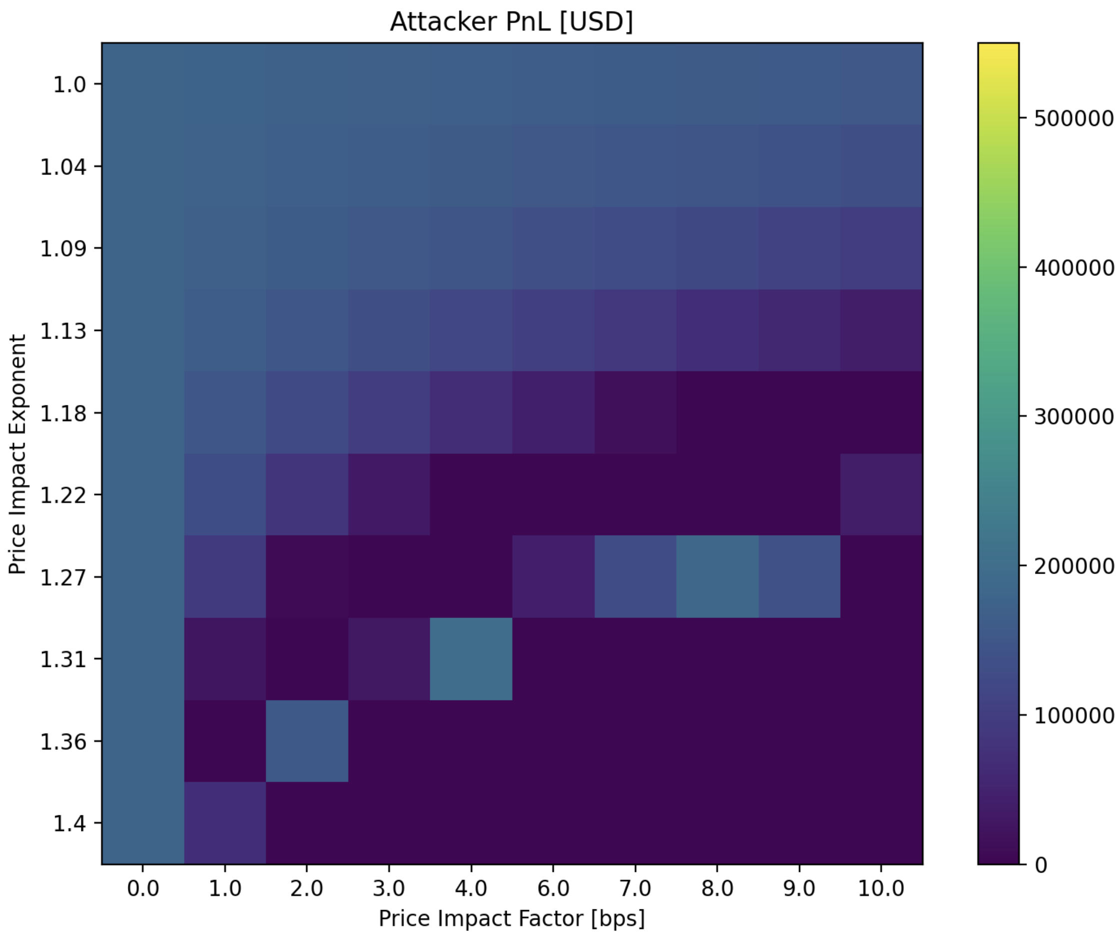

Figure 2.7: Our analysis evaluates the attacker's PnL across various price impact parameters. This investigation is conducted under a specific condition where the attacker employs a strategy of interleaving orders. For this scenario, the $minCollateralFactorForOpenInterestMultiplier$ is calibrated to $10^{-}7$.

3.8 Historical Implications for V1 Traders

Under the selected configuration for price impact parameters, we venture to examine its effects on 'typical' trading activities. To carry out this analysis, we deploy the data from GMX V1 positions associated with the asset under study to simulate the imbalance in an isolated V2 pool. Our primary aim is to quantify the fees users would have been compelled to pay due to price impact.

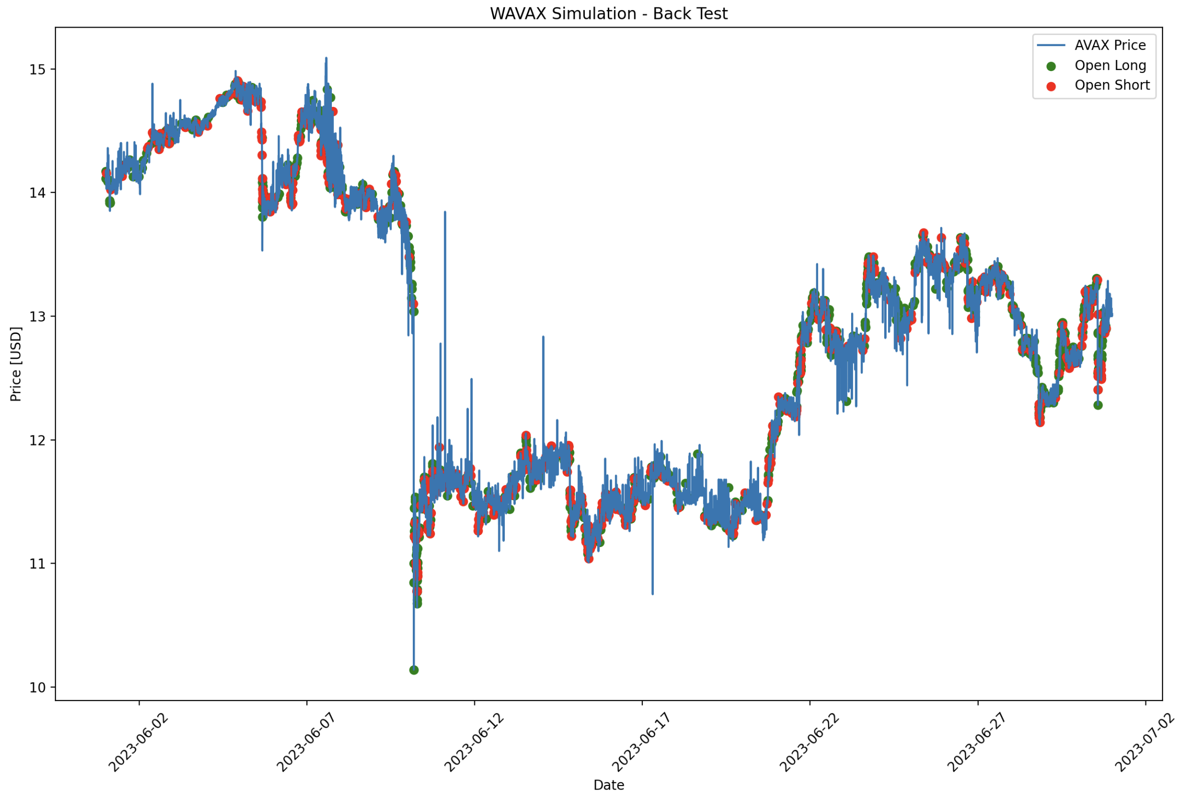

Figure 2.8: Time series of asset price and positions open times over 30 days. Note that closing times appear on these plots, but closing positions and realized PnL is also considered when tracking the pool imbalance and calculating price impact fees.

We can utilize a histogram depicting both negative and positive price impacts, with each observation corresponding to a completed trade wherein the taker experienced the depicted price impact. Such analytical, graphical representations provide profound insights into the distribution of price impacts and assist in comprehending the dynamics of trading activity and its consequent effects on market conditions.

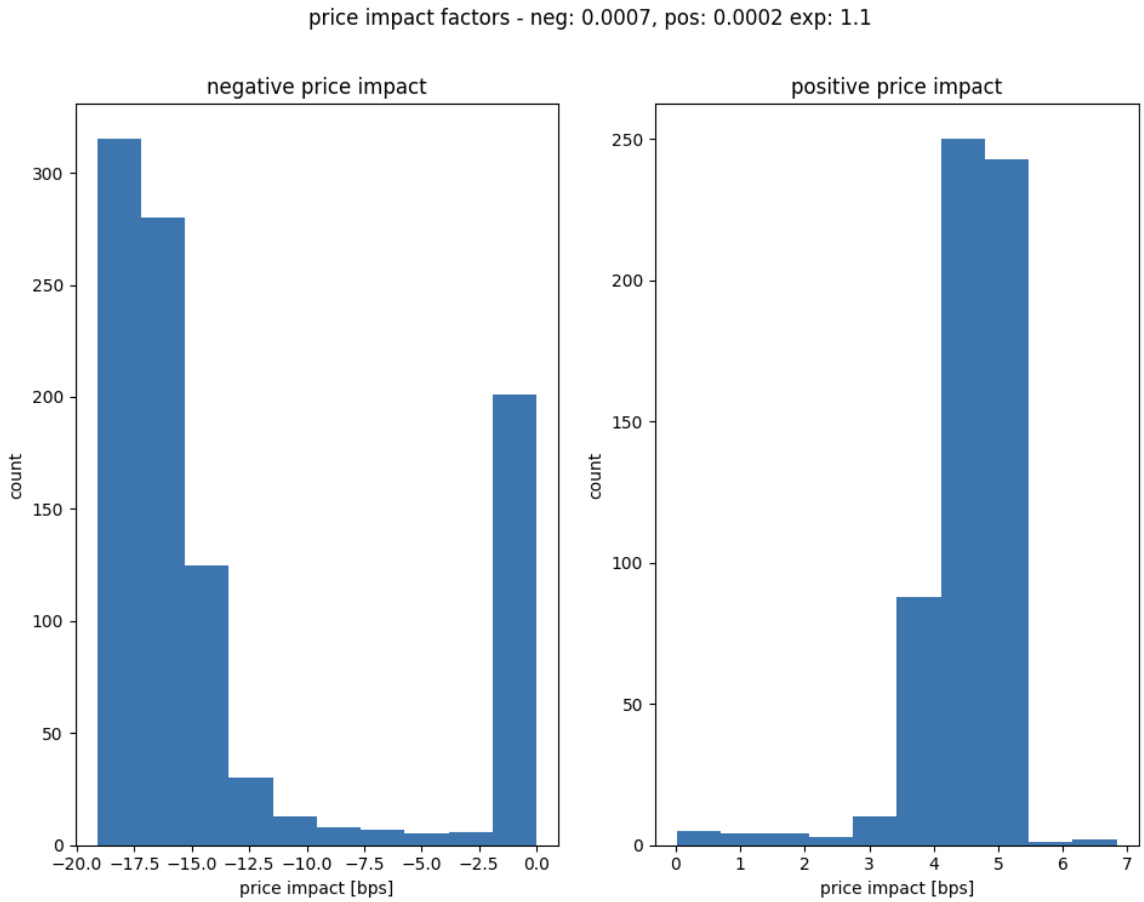

Figure 2.9: Time series of asset price and positions open times over 30 days. Note that closing times appear on these plots, but closing positions and realized PnL is also considered when tracking the pool imbalance and calculating price impact fees.

The frequency and magnitude of these price impacts can be utilized to assess the liquidity and volatility characteristics of the market, inform strategic decision-making processes, and assist in the calibration of trading models.

Considering the pronounced skewness of the pool for a substantial portion of the time, we observe that most of the 'destabilizing' trades were executed with a negative price impact of around 15-20 basis points. Along with position fees of approximately around 5 basis points, users would face a cumulative cost of less than 25 basis points when establishing positions.

This retrospective probe into the fees linked with destabilizing trades under the V2 model offers invaluable insights into potential changes in trading patterns and cost structures under diverse market conditions. It serves as a significant comparison of the trading landscape and could influence future trading tactics and risk management approaches.

3.9 Medium Market Cap Synthetics

While the methodology in this category bears substantial similarities to the preceding category, the profit and loss (PnL) are not constrained by pool size if the index asset appreciates more than the pool collateral. We set the reserve factor for the genesis pool at 50%, which will be revised based on historical price trajectories between the index token and

$longCollateralToken$.

4 Genesis Open Interest Cap Recommendations

We put forth recommendations for setting the initial open interest (OI) caps to ensure a secure launch, with these recommendations being informed by historical data from Version 1 (V1) of our platform. The aim is to strike an equilibrium between safety measures and robust user engagement while facilitating significant migration from V1. Therefore, we advise setting the initial caps at twice the mean open interest observed over the past 90 days on V1.

The lack of historical data mandates a distinct approach for tokens not listed on V1 (for instance, ARB). We suggest an initial cap of $1 million in these instances. This conservative estimate both safeguards the platform and encourages user engagement and growth. As we gather more data on these new tokens, the cap can be dynamically adjusted to mirror their individual market dynamics and risk profiles.

Asset	Long
WBTC	64M
WETH	64M
LINK	2M
ARB	1M
UNI	1M

Table 2.1: Arbitrum Open Interest Cap Recommendations

Asset	Long
BTC.b	64M
WETH	64M
WAVAX	2M

Table 2.2: Avalanche Open Interest Cap Recommendations

We emphasize that the proposed Genesis Open Interest Caps are chiefly designed to manage economic risk and do not explicitly account for the smart contract risk associated with freshly deployed contracts. If the GMX community prefers a more cautious strategy during the initial weeks post-launch, they may contemplate further lowering the caps. This prudent strategy would allow comprehensive monitoring of trading activity and empirical validation of system integrity before any cap escalation. This approach would be particularly advantageous during the initiation of the bounty program, which, despite its benefits, could also introduce additional uncertainties and risks. This methodology facilitates a safer exploration of the early dynamics of the platform, thereby aiding in the gradual establishment of a robust, resilient, and reliable system.

5 Funding Fees

Funding fees serve a pivotal role in encouraging a balance between long and short positions. The faction with a larger open interest pays a funding fee to its counterpart, which holds a smaller open interest.

The funding fee for the larger party is calculated using the formula:

$$MajorityFundingFee = \frac{fundingFactorPerSecond * openInterestImbalance^{fundingExponentFactor}}{totalOpenInterest} \tag{2.6}$$

The exploration of a dual-slope model, similar to the one used in Aave/Compound for interest rate calculations, warrants consideration. This model employs two distinct rates based on the pool's utilization rate, thus enabling the mitigation of lending and borrowing imbalances. A higher interest rate is activated when the utilization rate exceeds a certain threshold, and a lower rate is applied when the utilization rate is below this threshold.

The justification for choosing an exponent based on the USD imbalance rather than the ratio is to amplify the impact of large imbalances. This approach inflicts a more significant penalty for substantial deviations compared to minor ones.

What constitutes an acceptable imbalance depends on market-specific dynamics and the risk tolerance of the protocol. The acceptable imbalance can be adjusted in accordance with market conditions and risk management principles. The basis derived from Centralized Exchanges (CEX) can be used as a useful benchmark in this regard, representing the cost of maintaining a position over a long period. This CEX basis could be converted into an open interest (OI) imbalance by using it as a metric to determine the allowable discrepancy in long vs. short OI within the pool.

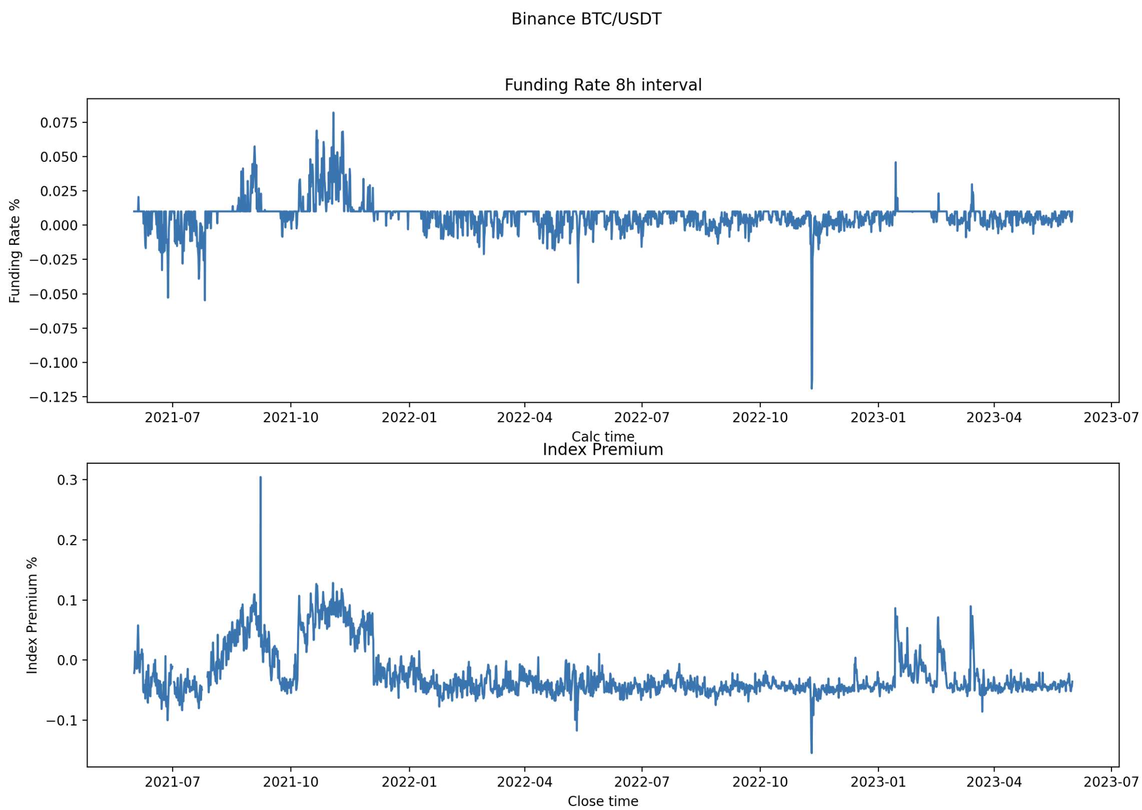

Figure 2.10: The visual representations above exhibit the Funding Rates and Index Premiums associated with Binance's most liquid perpetual market, BTC/USDT. The fees on Binance are calculated and collected over an 8-hour window timeframe.

The Premium Index function is formally expressed as follows:

$$PremiumIndex(P) = \frac{max(0, ImpactBidPrice - PriceIndex) - max(0, PriceIndex - ImpactAskPrice)^{fundingExponentFacto}}{PriceIndex} \quad (2.7)$$

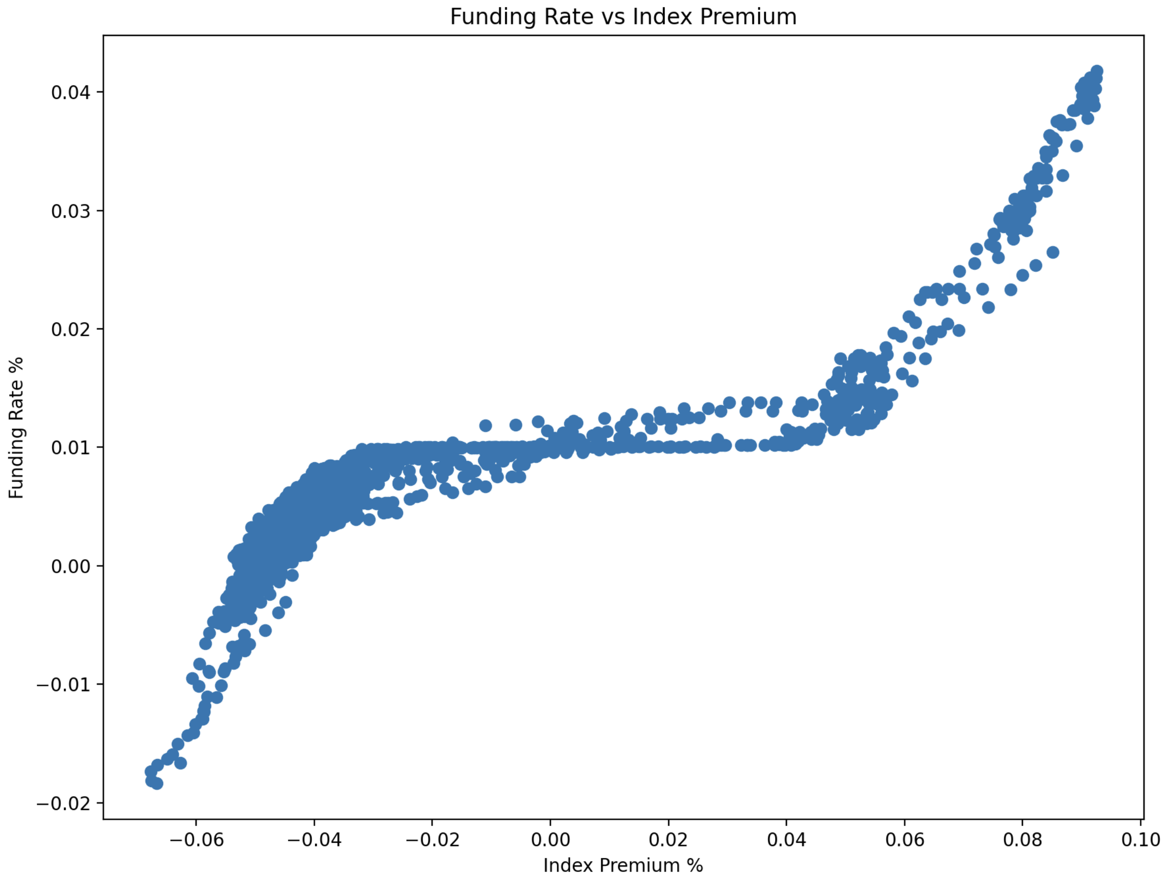

Figure 2.11: The interplay between the funding rate and the index premium should mirror the relationship between the funding rate and the pool imbalance, as observed on the GMX Synthetics simulations.

Dependent on the open interest skew, the operational principle of charging and rebating price impact mirrors the price difference observed between spot and futures prices. This is a gap that the Centralized Exchange (CEX) funding rates aim to reduce and eventually eliminate.

The market dynamics at play are the same; the primary difference lies in the role of market makers. Within GMX, Liquidity Providers (LPs) serve as the market makers absorbing excess exposure until traders intervene to restore the market's midpoint, as opposed to this function being fulfilled by traders or market makers in a conventional CEX scenario.

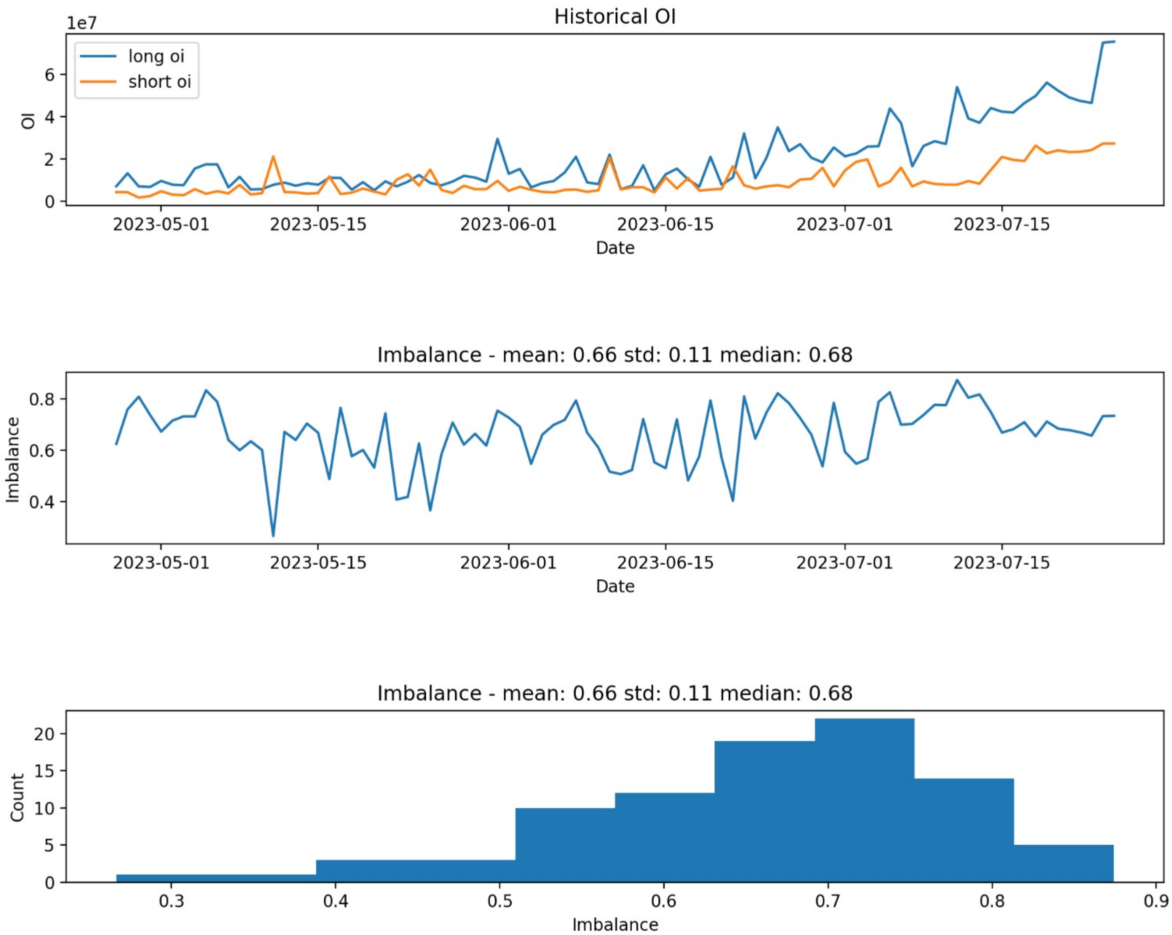

Figure 2.12: The figures presented herein offer a multifaceted interpretation of Historical Open Interest, defined as the proportion between long and short positions within the GMX Liquidity Pool (GLP). The three visualizations each underscore unique aspects of this parameter.
The first illustration plots a temporal series of Open Interest imbalances, effectively illuminating the directional skew of the pool over time. This chart provides a dynamic view of how the pool's inclination towards long or short positions fluctuates with market conditions.
The subsequent figure narrates a time-series representation of open interest imbalances. The notable observation is the mean skew of 74:26 and a standard deviation of 7%. This depiction indicates a consistent and significant tilt toward long positions throughout the examined period, with a degree of variation encapsulated by the standard deviation.
The final chart offers a histogram of historical open interest imbalances. This portrayal presents a bucketed measure of pool skew, effectively summarizing the distribution of the imbalance ratios. The median skew demonstrated is 77:23, indicating that long positions often outnumbered short positions, revealing a consistent pattern of trader sentiment within the pool.

The above graphics provide a historical overview of the GMX V1 model's imbalance. By studying these historical imbalances, we can empirically understand what the market considers an acceptable level of pool skew. While we do observe anomalies during periods of heightened volatility (June 10th, 2023 is a notable example, showcasing the fallout from SEC comments on major crypto assets being securities, thus precipitating a market-wide sell-off), a consistent demand curve emerges, most frequently indicating a skew towards longs in an 80:20 ratio, with a standard deviation of 0.7.

Armed with this information, our objective becomes the optimization of the interplay between funding rates and open interest (OI) imbalance, concentrating our efforts on the [70,30]-[80,20] ranges. When the skew veers towards a balanced 50:50 ratio, the associated payments become negligible due to the minor imbalance. On the other hand, when the skew rises above this level, it triggers more substantial funding fees, thereby incentivizing traders to strive towards reestablishing equilibrium within the pool.

5.1 Analyzing Genesis Funding Fee Parameters

Option 1: Exponential Model with Funding Exponent >> 1

Utilizing an exponent value greater than one introduces a high incentive to adopt the minority side while penalizing the majority side heavily. This, however, creates a dependency on the size of the open interest (OI) rather than solely on the imbalance. While this might be acceptable under the assumption of a "slowly changing" OI size, as has been observed historically, we scrutinize the varied funding rate curves per OI size under different scenarios.

Example 1.a: Total OI = 4M , Exponent = 1.8, Funding Factor = 2.5e-13

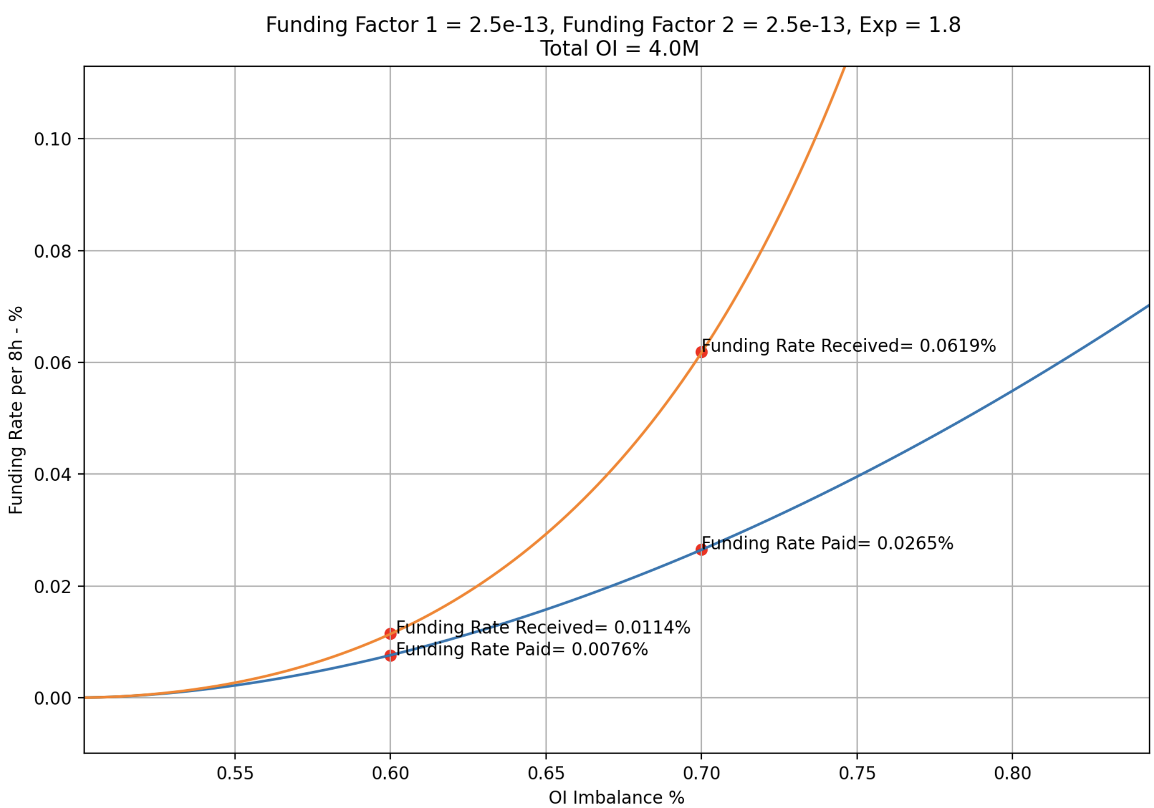

Example 1.b: Total OI = 50M , Exponent = 1.8, Funding Factor = 4e-14

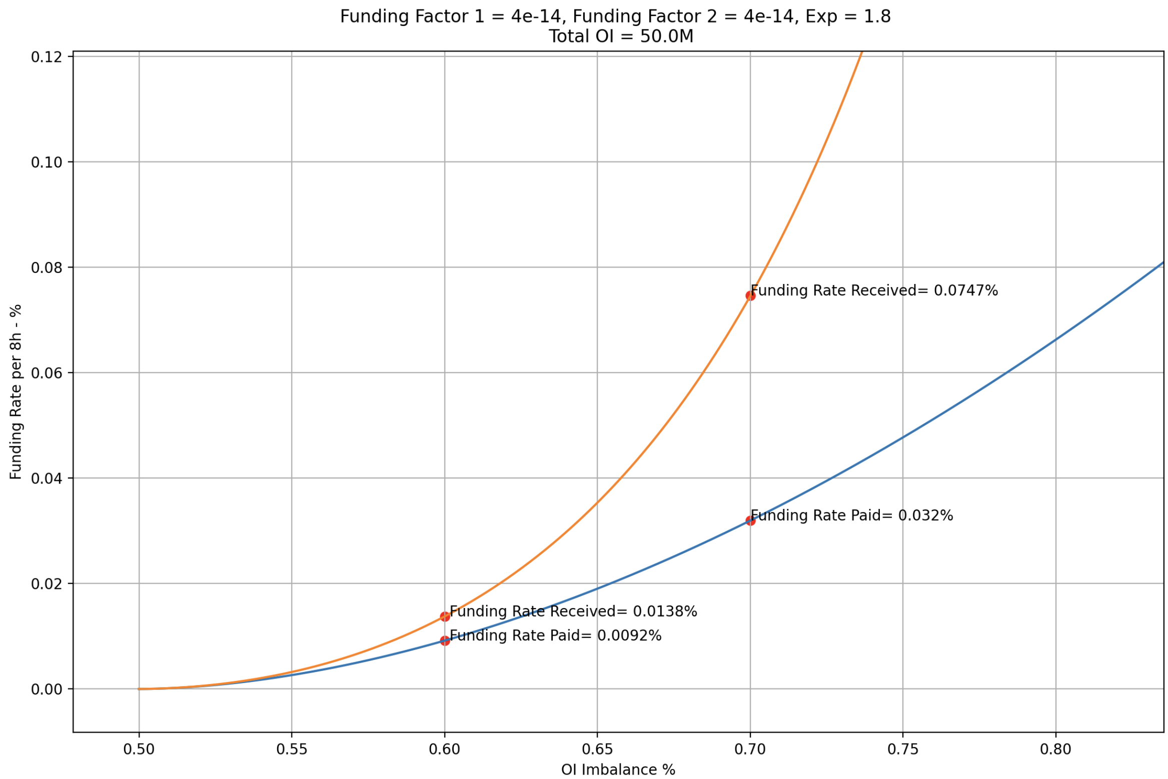

Example 1.c: Total OI = 100M , Exponent = 1.8, Funding Factor = 2e-14

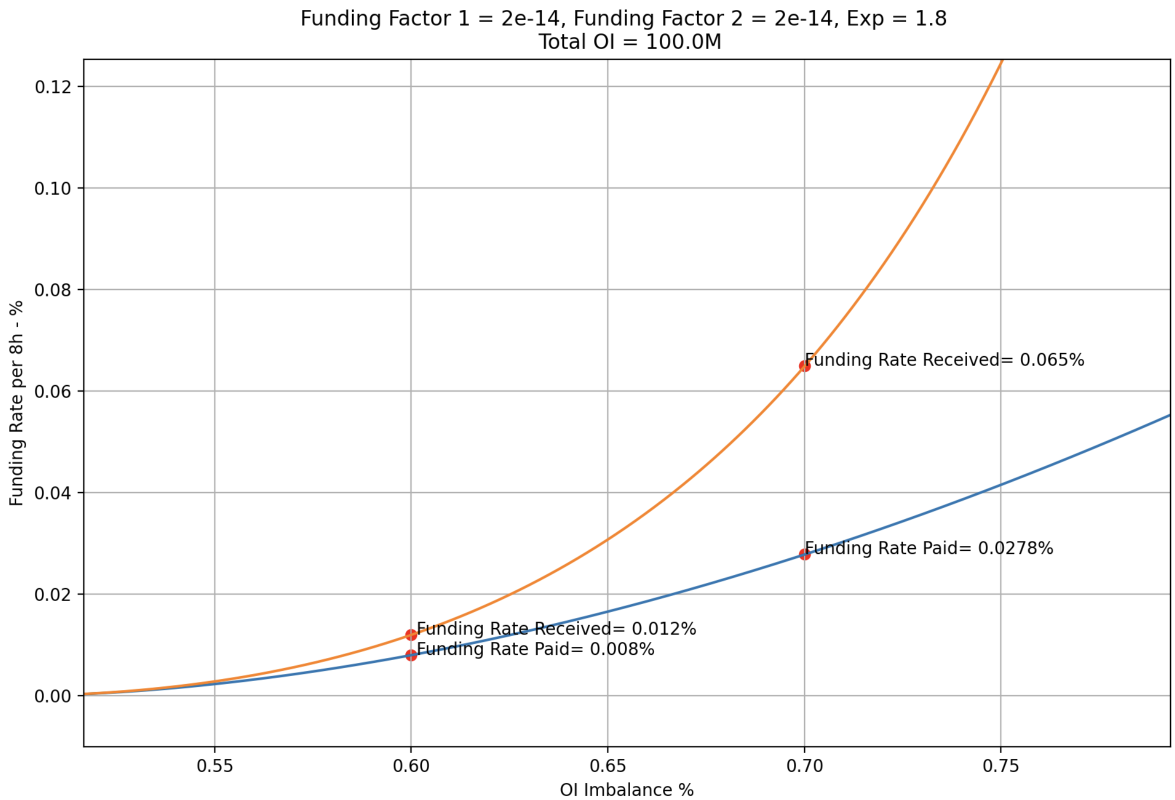

Option 2: Linear Model with Funding Exponent = 1, Factor = $5e-8$

Implementing a linear model eliminates the dependency on OI size and proposes a significant incentive to adopt the majority side. Nevertheless, the funding rate paid by the majority grows linearly, which may not adequately disincentivize taking the majority side. This could result in a balance failure or an excessively high utilization rate.

Example 2.a: Total OI = 1M , Exponent = 1, Funding Factor = 2e-8

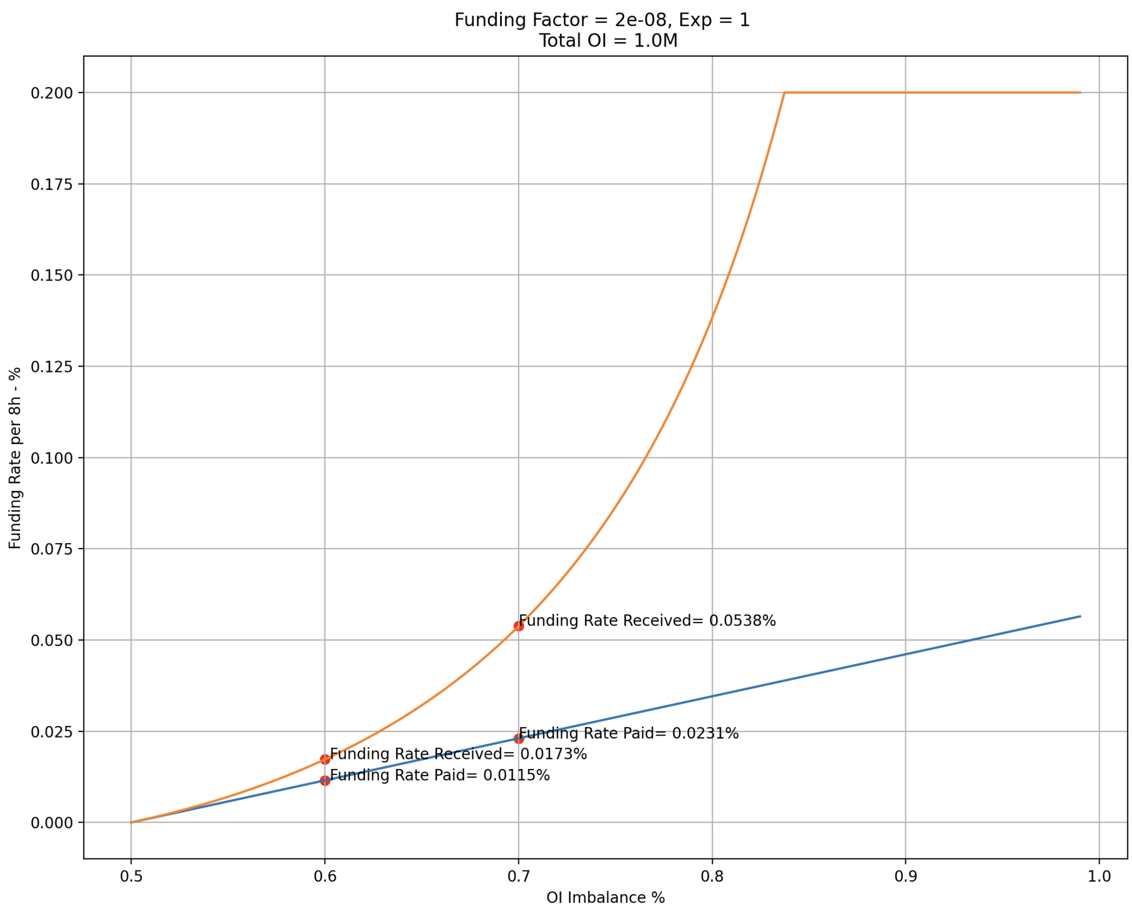

Option 3 (Recommended): Kink Model with Two Linear Phases

This alternative successfully removes the dependency on OI size while continuing to be an effective rebalancing mechanism. However, it requires dynamic altering of the $priceImpactFactor$ when the skew reaches a certain threshold (70%).

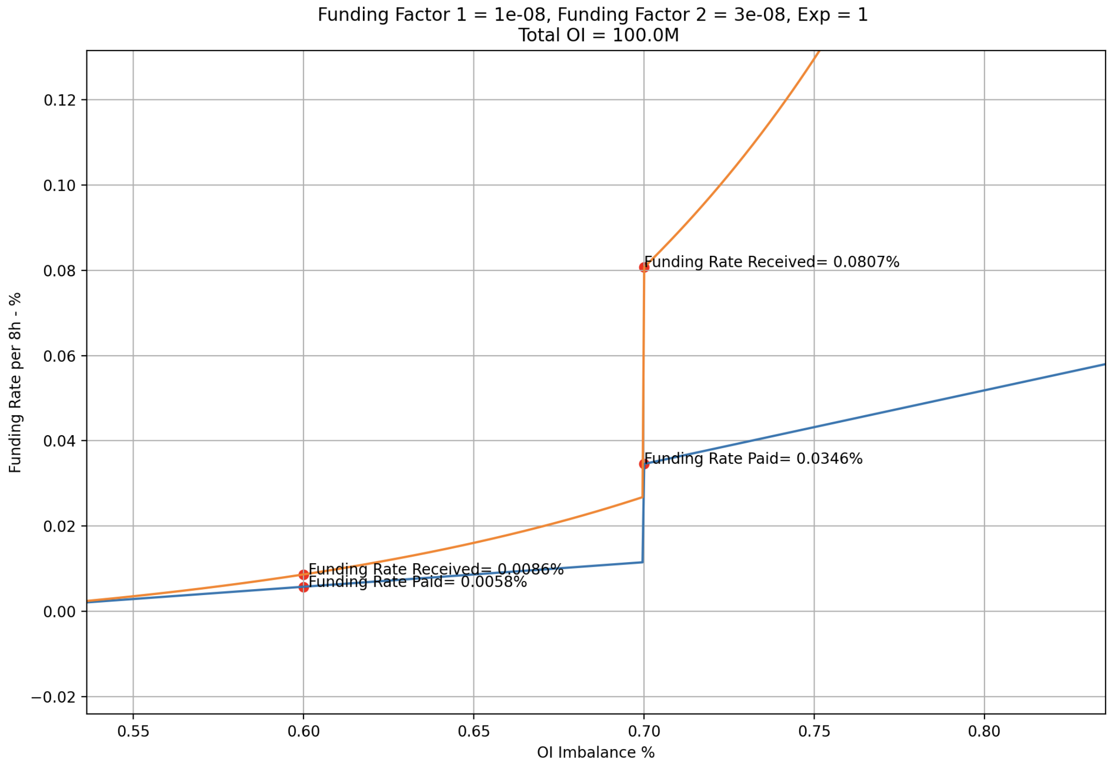

6 Borrowing Fees

A borrowing fee is paid to liquidity providers; this helps prevent users from opening long and short positions to take up pool capacity without paying any fees.

Borrowing fees are calculated as follows across longs:

$$LongBorrowingFees = \frac{borrowingFactor * (OIinUSD + pendingPNL)^{borrowingExponentFactor}}{poolFundsInUSD} \tag{2.8}$$

Borrowing fees are calculated as follows across shorts:

$$ShortBorrowingFees = \frac{borrowingFactor * (OIinUSD + pendingPNL)^{borrowingExponentFactor}}{poolFundsInUSD} \tag{2.9}$$

For example, let's examine the following scenario:

- The borrowing factor per second is $\frac{1}{50{,}000}$
- The borrowing exponent factor is 1
- Long open interest is \$150,000
- Pending PnL is +\$50,000
- The pool has \$250,000 worth of tokens.

The borrowing fee per second for longs would be calculated as follows:

$$(1/50,000) * (150,000 + 50,000)/250,000 = 0.000016\% = 0.0016\% \tag{2.10}$$

6.1 Utilization Rates

Utilization rates are expressed as:

$$utilization \; = \; \frac{OI}{poolSize} \tag{2.11}$$

When it comes to utilization rates, our strategic goal varies depending on the existing level of utilization. At lower utilization levels, our aim is to stimulate more trading activity and encourage the opening of additional positions. This helps to optimize the use of the pool's resources and promotes a vibrant and active market environment.

On the other hand, when utilization is extremely high (approaching 100%), our strategy shifts to discouraging traders from maintaining their open positions. This step is necessary to ensure that liquidity providers are able to withdraw their capital and prevent a problematic lock-up of resources. This approach maintains the pool's flexibility, ensuring its sustainability and resilience under a range of market conditions.

6.2 Interplay Between Borrowing Fee and Funding Fee: An Analytical Perspective

The mutual influence of borrowing and funding fees greatly impacts the effectiveness of Open Interest (OI) balancing mechanisms. To understand this complex relationship better, we can consider the following theoretical scenario:

- The OI is skewed 40/60 in favor of longs/shorts.
- The utilization is measured at 30%.
- Holders of long positions are subject to a funding fee of 5%, while holders of short positions receive a return of 7.5%.

In these conditions, assuming a maximum borrowing fee of 10% with an exponent of 1, the effective borrowing fee comes out to 3%. After factoring in the borrowing and funding fees, the net outcome for short positions is a return of -4.5%.

However, this equation becomes more complex when the exponent is adjusted to 2 or 3. Such a change could reduce the borrowing rate for low to medium utilization levels, possibly resulting in a significantly positive funding fee.

This simulation highlights the intricate and non-linear relationship between borrowing and funding fees. Understanding these dynamics is crucial when developing financial models. Thorough and rigorous analysis is necessary to balance capital efficiency and risk mitigation, emphasizing the importance of optimally calibrating these parameters. Since we don't have data on user elasticity to the new fee structure, we eliminate borrowing fees from the minority side on launch, to assure a high incentive to rebalance the pools.

6.3 Methodology Constraints And Relaxations

Given that V2 remains in its pre-launch phase, our investigative approach predominantly relies on a set of conjectures owing to the absence of empirical data. For instance, we have incorporated agents, the conduct and risk proclivities of whom are extrapolated from those observed in V1. Furthermore, we have synthesized price trajectories and liquidity landscapes predicated upon historical market price oscillations. While we view V2 and V1 as distinct financial instruments, we expect that the behavior empirically witnessed will likely deviate due to factors such as the establishment of isolated pools and the genesis of novel pools.

The generation of synthetic data predicated on historical observations, although potent, does possess inherent limitations as it posits that all future market movements or scenarios are contingent upon historically discerned conditions. Given the recurrence of unanticipated 'Black Swan' incidents within the crypto realm, an extension in the operational period of the system would augment the data set we can accrue, thereby enhancing the accuracy of our models.

In light of these considerations, our methodology and initial genesis suggestions are provisional and mandate successive iterations as more data surfaces. At Chaos Labs, we place the highest emphasis on a secure and methodical launch, prompting us to adopt a prudent stance regarding specific modeling parameters.

For example, within the scenario of the interleaving order attack, we presuppose that an attacker, gauging the potential profitability of the assault, would aspire to a twofold return on her capital expenditure. This assumption could potentially overstate the risk tolerance of some attackers who might be inclined to jeopardize capital and instigate an attack with considerably lower potential yields.

We maintain adaptability with these parameter recommendations, and modifications can be enacted per the community's risk appetite. Accommodating these alterations, we can conduct new simulations to yield an updated set of genesis parameters, refining our model to better mirror market realities and potential risks.

7 Future Model Improvements

The capabilities of the Chaos Labs platform are being amplified to amass on-chain user behavior data and patterns pertinent to the utilization of V2. The initial weeks following the V2 launch are projected to offer opportunities for data gathering and processing, enabling further refinement of our models, proposing adjusted risk parameters, and gauging user elasticity in response to various market fees. This period will necessitate an additional research sprint and a new cycle of simulations. We are devoted to facilitating this rigorous analytical phase to assure a secure and optimized protocol environment.

Chapter 3

Agent-Based Modeling and Simulations

1 Simulation Framework

In our research, we construct advanced simulations that mirror the operations of on-chain protocols. These simulations rely on mathematical premises grounded in relationships between protocol parameters and external variables. Collectively, these assumptions form a model, serving as an analytical tool to understand the functionality of the corresponding protocol more profoundly.

Simulations become the instrument of choice when the relationships that constitute the model are complex. In instances where the relationships are linear, it might be plausible to employ closed-form mathematical techniques or equations for obtaining precise results, a practice referred to as analytical solutions. However, the overwhelming intricacy of most real-world systems makes analytical models a nonviable solution, necessitating the need for simulations. In a simulation, we employ computational means to assess a model numerically to approximate the model's desired actual attributes.

Yet, despite their value, simulations are not ubiquitously used. The process of creating high-precision simulations is complex and laborious. The task is particularly challenging when modeling DeFi protocols, given the nascent nature of this field with limited historical data and relatively small datasets. The challenges primarily fall into two categories: precise modeling and scalability. In the subsequent sections of this paper, we outline the system architecture we have designed to address these unique scalability challenges and focus on the simulation modeling and statistical framework used to generate recommendations.

Parallel simulations and high-level system architecture are critical components of Chaos Labs' approach. Chaos Labs features a custom EVM Simulation environment based on a proprietary, Python-constructed EVM. These simulations initiate with a data-synchronization phase, during which new or historical mainnet data is extracted and loaded. This data spans various dimensions, including account portfolios and balances, agent elasticity, protocol liquidity, and risk parameters. The performance of Chaos EVM simulations is noteworthy, displaying a 250-fold improvement in latency and CPU performance compared to their on-chain counterparts.

DeFi applications execute on data-rich, transparent blockchain environments. This level of visibility significantly surpasses that of traditional finance applications, which operate in more opaque settings. The capacity to fork blockchains delivers inherent transparency and data readability. On-chain simulations are conducted by a simulation executor utilizing a novel agent and scenario model to interact with a dedicated blockchain fork. This unique approach allows us to obtain a comprehensive snapshot of the runtime environment at any given block height, while agents can interface with blockchain protocols identically to how they would interact in a production environment.

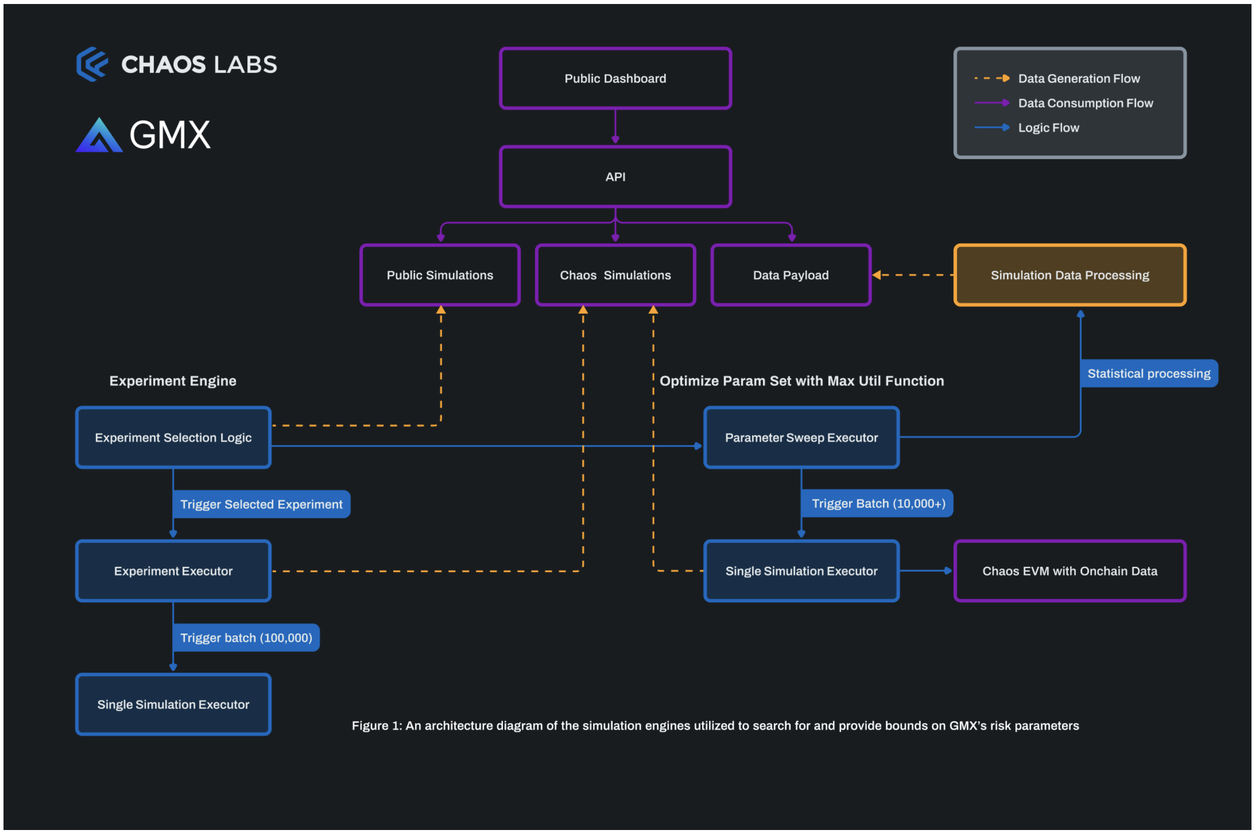

Figure 3.1: High-level diagram of the Chaos Labs Simulation Architecture

1.1 Implementing Tools for High-Fidelity Simulation Analysis: A Comprehensive Evaluation

The unique Chaos Cloud Architecture addresses the large-scale requirements of value-at-risk estimation. Our proprietary cloud solution parallelizes thousands of executors and forks across a multitude of machines while real-time data processing enables rapid statistical analysis. The process begins with the initial parameter exploration executing on the proprietary Chaos EVM, significantly reducing the search space and producing a parameter recommendation. This recommendation is subsequently backtested, and the results are shared within the community to uphold verifiability and transparency. An illustration of this architecture can be seen in Figure 3.1.

Agent-Based Simulation (ABS) involves the employment of agents as autonomous entities capable of perceiving their environment, including other agents, and utilizing this information for decision-making processes. In the context of DeFi simulations, agents mirror the actions of users, encompassing traders, arbitrageurs, liquidators, and exploiters, among others. This ABS allows for dynamic interaction and adaptation within the environment, as demonstrated in the GMX example, where traders, liquidators, and arbitrageurs interact with GMX and DEXes in response to price variations.

Monte Carlo Simulations introduce a degree of randomness by assigning a random value to the uncertain variable, calculating the model output, and repeating the process with different variable values. The results of multiple iterations are aggregated and averaged to obtain an estimate. For instance, in the case of GMX, each simulation applies a randomly generated price trajectory, using an agent-based model to simulate the response this elicits from borrowers and liquidators of the protocol.

Agent-Based Monte Carlo Simulations are particularly relevant for financial applications, where price dynamics fully describe a static portfolio's future value. In contrast, our problem is to simulate potential protocol losses, a measure contingent on price dynamics and the interactions of traders, exploiters, arbitrageurs, and liquidators. Hence, we update the actions that traders and arbitrageurs would implement for each simulation at that specific timestep. As long as our simulations are statistically unbiased, we can utilize these simulation results to estimate the expected value of protocol losses. Further in the paper, we elaborate on the methodology used to construct these Monte Carlo agent-based simulations, and the statistical testing framework employed to determine the value at risk for specific parameter configurations.

2 Taxonomy of Agents in Agent-Based Simulations

In an agent-based simulation, the role of the agents is instrumental. They represent individual entities within the system that have the ability to act autonomously, sense their environment, and adapt their behaviors based on the inputs and the interactions with other agents. This taxonomy, therefore, provides a structure for understanding the different types of agents and their respective roles within the simulation. We delineate the agent taxonomy into three high-level groups:

1. **Malicious actors**: These agents strategize to exploit periods of low liquidity to manipulate prices, deploying limit orders integrated with stop losses and take profits.
2. **Arbitrageurs**: These agents are engaged in arbitrage between the spot market and the perpetual futures markets within the same venue and across Centralized Exchanges (CEXs).
3. **GMX Traders**: These are benign actors operating on GMX-Synthetics. The distribution of their positions is extrapolated from GMX-V1 historical data and subsequently adjusted to align with a variety of potential growth scenarios.

Further, we detail the different flavors of agents across each type:

Table 3.1: Perpetual Trader Taxonomy

Agent Role	Agent Model/Type	Description
Perp Traders	Trend Traders	Traders that trade according to the prevailing trend in the market.
	Short Term Hedging Traders	Traders use GMX to hedge positions opened elsewhere or spot positions in their wallets. These positions are closed after a short period once the exposure to the underlying asset is closed.
	Long Term Hedging Traders	Similar to short-term hedging traders, these traders maintain positions elsewhere over time (typically long) and hedge them to be delta-neutral or reduce exposure to a predefined level.
	Volatility Traders	These traders don't care about the direction of the price but rather trade the volatility of an asset. They typically open long and highly leveraged short positions of similar sizes when they expect very high volatility.
	On-chain/Off-chain Perpetual arbitrageurs	Will buy or sell any amount as long as the counter trade can be executed on a CEX with significant margin to account for the time risk between execution on GMX and CEXes.
	Funding traders	Similar to arbitrageurs, these traders will hold opposite positions on GMX and another exchange as long as the net funding is positive and outweighs other fees.
	Price Impact Traders	Opening opposite positions on GMX and a CEX when the price impact on GMX is so high that the net fees are positive once the pool is balanced again.

Table 3.2: Liquidity Provider Taxonomy

Agent Role	Agent Model/Type	Description
LPs	Static LPs	LPs that trust the protocol and provide long-term liquidity. They add or remove liquidity based on the long-term yield that LPs earn.
	Opportunistic LPs	LPs that identify specific market conditions (like low volatility) or pool states that yield excess yield. They only add liquidity when terms are extremely favorable and pull out when it's less favorable.

Table 3.3: Spot Trader Taxonomy

Agent Role	Agent Model/Type	Description
Spot Traders	Trend Traders	Selling or buying specific tokens when there is excess supply or demand in the market for the token. They are less price sensitive than fully rational actors.
	On-chain spot arbitrageurs	Will buy or sell any amount as long as the counter trade can be executed atomically on another DEX and yield profit.

Table 3.4: Malicious Agent Taxonomy

Agent Role	Agent Model/Type	Description
Malicious Actors	Market Manipulators	Whales that can move CeFi markets and manipulate the Oracle price.
	Market Makers	CeFi Market Makers can manipulate markets by draining liquidity for short periods to create increased volatility

Following the collation of data from production, future research will be directed toward the integration of more sophisticated agents:

1. **Basis traders**: These agents capitalize on high projected funding fees by taking the minority side.
2. **Positive price impact fees arbitrageurs**: These traders adopt strategies to compress the skew by opening positions, thereby earning price impact fees. Such traders might opt to hedge their positions on other venues to manage risk.
3. **Hedging agents**: These are DeFi Liquidity providers/Arbitrageurs who hedge their exposure on GMX - Synthetics.

3 Computing price trajectories

We use discrete-time random walk stochastic processes to generate price trajectories of assets.

Geometric Brownian Motion. The canonical model for generating price trajectories is geometric Brownian motion (GBM). In a simple GBM model, the price (S) changes according to a stochastic differential equation: $dS_t = S_t(1 + \mu dt + \sigma dW_t)$, where W_t is a Brownian motion, σ is a volatility scaling parameter, and μ is a forcing parameter. This stochastic DE can be integrated with Itô's formula to get that $S_t = S_0 \cdot e^{(\mu - \frac{\sigma^2}{2})t + \sigma W_t}$. This analytic solution makes it possible to sample a price trajectory's outcome without running each timestep.

Issues with GBM. Although GBM is a popular model in finance literature, it does not hold in practice at short intervals. We often see erratic price jump behavior at short timescales and short bursts of high or low volatility, which the constant-σ GBM model does not reflect. We also see times of new information that lead to a short-term jump or drop in

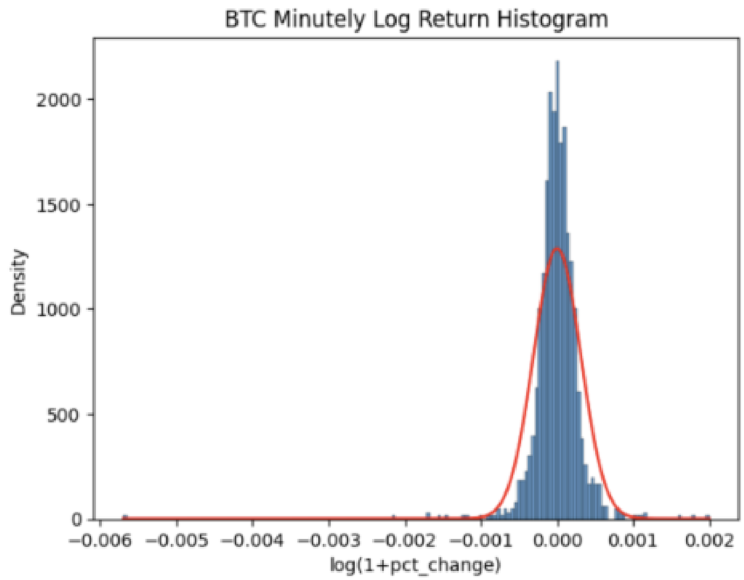

Figure 3.2: A histogram of BTC minutely log returns from 11/19/2022 to 11/20/2022. The red line is a normal distribution with the same mean and standard deviation as the log returns. That is, BTC minutely returns are lognormal.

returns, which is also not reflected in the constant-σ GBM model. To correct for these, we forego the GBM modeling and instead utilize variable volatility price trajectories.

For clear evidence of the non-normality of this data, see the distribution of price trajectories in figure 3.2. We see that the distribution of returns contains outliers that are many standard deviation events. To model returns as following a GBM process, we would either need to entirely ignore long-tail events by using the real σ, or we would need to make these long-tail events possible by increasing the σ, which would overestimate the volatility of returns in the majority of cases. A GBM with constant σ is not a suitable model for returns behavior, and this is of particular importance because long-tail returns have a significant impact on the health of GMX accounts.

Not only is the returns distribution fat-tailed, but it is also autoregressive. Figure 3.3 shows a plot of the minute log returns for over 1 month, during which there were several high-volatility events. In some of these periods, the median across 100 minutes also elevates. This demonstrates that a minute with a significant absolute return is likely to be close to other high absolute return minutes. Assuming that the mean return is approximately zero, the returns are simply the residuals of a time series process with a mean of 0. The variance of these residuals demonstrates autoregressive tendencies.

Modelling Volatility with GARCH. We can extend our GBM model to track more realistic volatility by modeling the autoregressive tendency of the variance of returns. Looking at historical data, we can see that a GARCH(1,1) model of the volatilities is the most fitting. From here, we can fit a GARCH model to historical data to find the baseline volatility (ω), the ARCH parameter (α), and the GARCH parameter (β). We then compute the

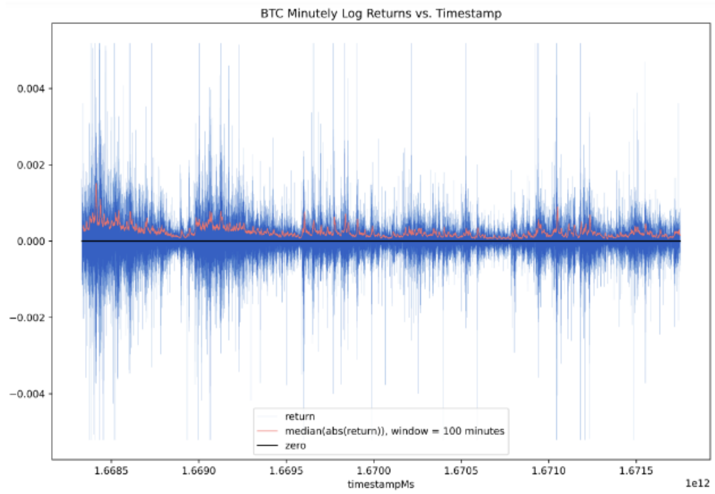

Figure 3.3: BTC log returns vs. time (blue) and BTC median of absolute log returns (red), from 11/14/2022 to 12/22/2022. The top and bottom 0.05% of values are removed for ease of viewing.

following (assuming that the drift, μ, is equal to zero):

$$S_t = S_{t-1}\left(1+\varepsilon_t\right), \tag{3.1}$$

$$\varepsilon_t = \sigma_t \cdot z, \text{ and} \tag{3.2}$$

$$\sigma_t = \sqrt{\omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2}. \tag{3.3}$$

The z_t term here is a white noise term with a mean of 0, and it is commonly set to z_t $\mathcal{N}(0,1)$. This new process is quite similar to the GBM model, except it is discrete, assumes zero drift, and has a time-varying volatility parameter σ. This process allows us to generate price paths for assets like we could for the GBM model.

Correlated GARCH. Although our GARCH model improves upon the GBM model, we do not yet capture the fact that returns are correlated. For GBM models, this is typically performed by requiring the white noise term for any two assets, i and j, their Wiener process terms, dW_t^i and dW_t^j, must satisfy $\mathbb{E}[dW_t^i \cdot dW_t^j] = \rho_{i,j}$, where $\rho_{i,j}$ is the Pearson correlation coefficient of the log returns. Instead of using the independent returns white noise distribution $z_t^i \sim \mathcal{N}(0,1)$, we sample all of the white noises from a multivariate normal distribution with mean $\mathbf{0}$ and covariance matrix $\Sigma = [\rho_{i,j}]$:

$$\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, [\rho_{i,j}]). \tag{3.4}$$

When assets' returns have no linear correlation, our formulation is identical to the white noise distribution for independent returns. In practice, many crypto assets are correlated. See figure 3.4 for a comparison of the multivariate normal distribution that we sample for ETH and BTC returns, compared to the returns that have arisen historically. See also figure **??** for an example of a single day's change in asset prices.

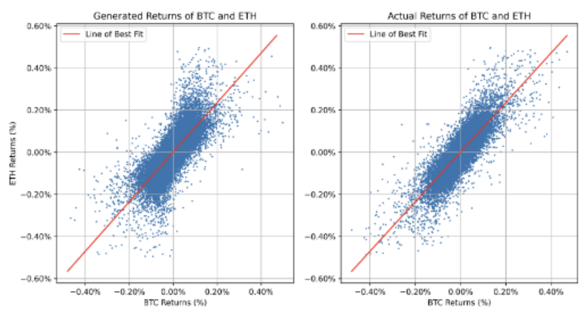

Figure 3.4: Left: returns generated by our multivariate white noise distribution for BTC and ETH. Right: historical returns for BTC and ETH. Each uses the same line of best fit, which was generated via historical data, is plotted on each.

Chapter 4

Future Risk Optimization and Robustness Analysis

1 Risk Optimization and Resilience Analysis

The precision of Agent-Based Modeling and Simulations hinges on the integrity and quality of the data employed as model input. The initial parameter recommendations are crafted based on a blend of security-focused economic heuristics, an examination of trader and liquidity provider (LP) behavior documented in Version 1 (V1), as well as an in-depth investigation of empirical trading patterns prevalent in centralized exchanges (CEX), supplemented by other empirical evidence. Collected data is normalized, tested, and verified against the blockchain, utilizing the Chaos Labs GMX V1 Risk Observability and Monitoring Hub.

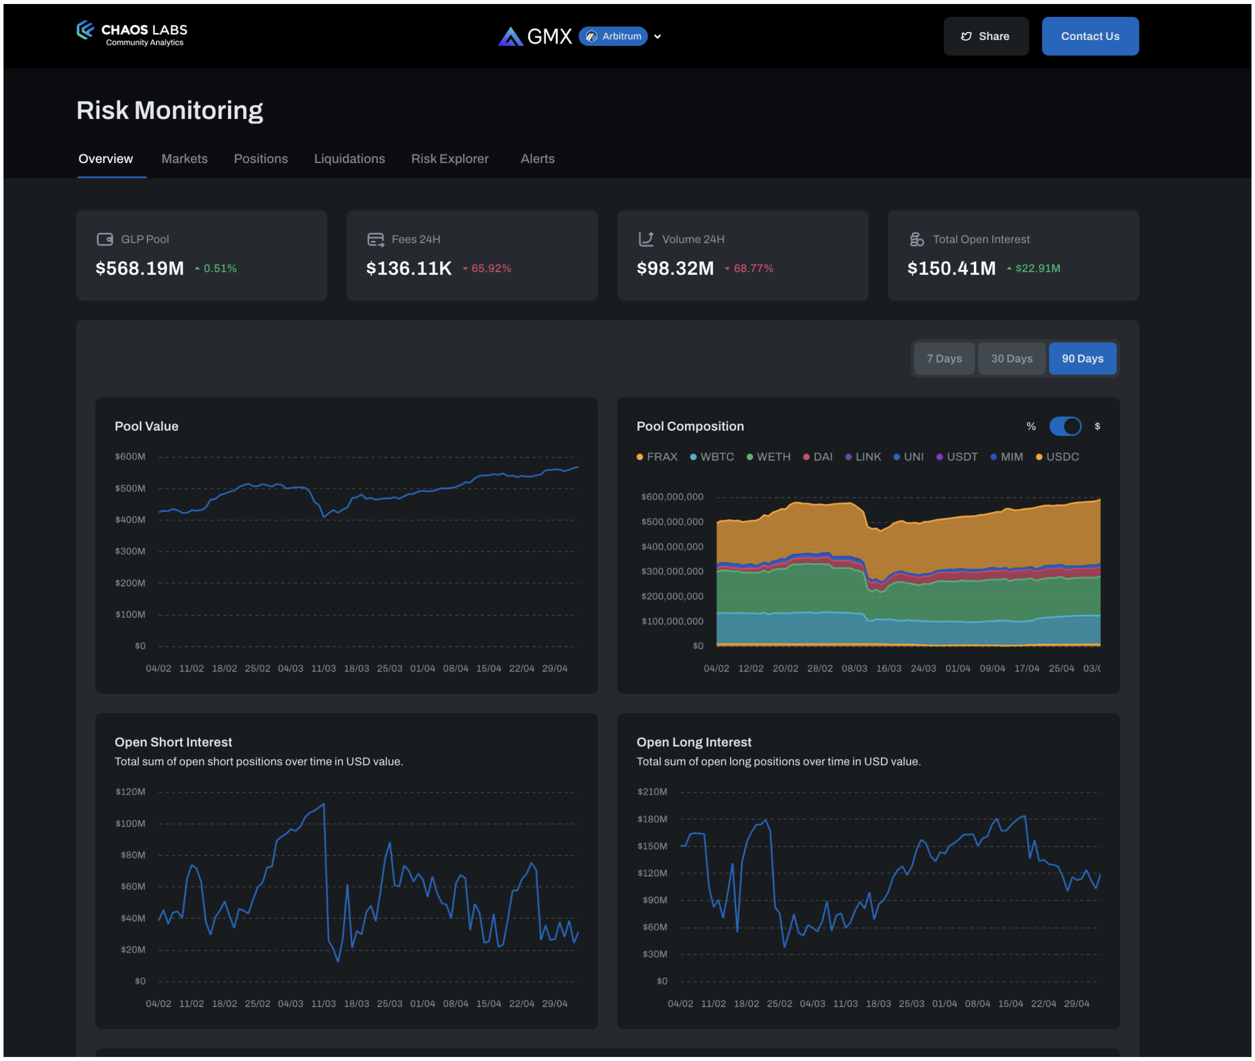

Figure 4.1: GMX Risk Hub streams real-time application data, and serves as a baseline data model for ABS simulations.

Despite the resemblances between Version 1 (V1) and Version 2 (V2), we interpret them as separate financial applications. As such, we foresee the emergence of unique agent behavior, trading strategies, and arbitrages, among other aspects. Post the V2 launch, Chaos Labs will maintain vigilant observation of platform utilization and incrementally refine our models, increasingly prioritizing the analysis of agent behavior within the Synthetics environment. As a consequence, these models and parameter recommendations will be subject to further refinement and optimization, guided by these empirical observations.

Appendix A

Medium Market Cap Asset Volatility

1 Observed Volatility in Billion Dollar+ Market Cap Assets

On June 6, 2023, the U.S. Securities and Exchange Commission (SEC) brought charges against Coinbase, Inc., alleging operation of an unregistered national securities exchange, broker, and clearing agency. The SEC also accused Coinbase of failing to register its crypto asset staking-as-a-service program. According to the Commission, Coinbase has been unlawfully facilitating crypto asset securities transactions since 2019, intertwining the traditional roles of an exchange, broker, and clearing agency without requisite registration, thereby denying investors of significant protections.

Furthermore, the SEC asserts that Coinbase has been offering an unregistered securities offering via its staking-as-a-service program since 2019. This program allowed customers to earn profits from blockchain transaction validation services, using their stakeable crypto assets. Coinbase's parent company, Coinbase Global Inc., is also deemed liable for certain violations, according to the SEC.

The SEC, voicing concerns over investor protection, seeks injunctive relief, disgorgement of ill-gotten gains with interest, penalties, and other equitable relief. The complaint also underscores the importance of compliance with federal securities laws and criticizes Coinbase's intentional disregard of these requirements. The SEC's charges have been filed in the U.S. District Court for the Southern District of New York.

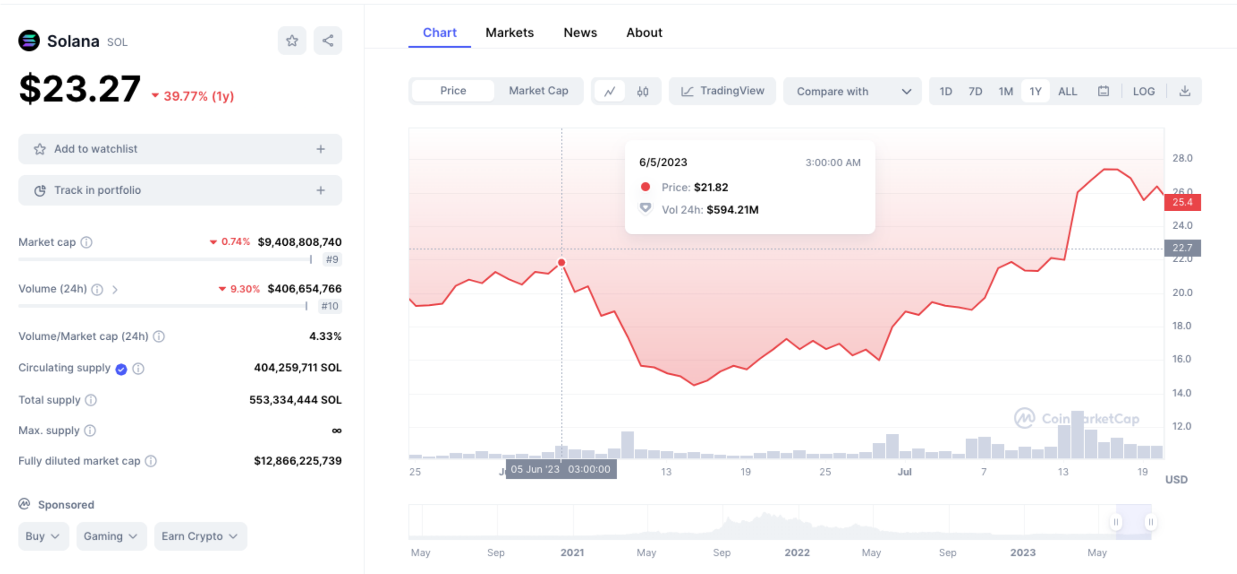

Figure A.1: This time-series price data chart, sourced from CoinmarketCap, displays the price action of Solana, a 5+ billion USD marketcap asset, 24 hours after launch.

The time series graph provided, delineates the price fluctuations of Solana, a digital asset with a market capitalization exceeding 5 billion USD, during the initial 24-hour subsequent to the SEC filing.

The inherent volatility of cryptocurrency markets is prominently exhibited in this scenario, as Solana undergoes a precipitous contraction in market capitalization of 33%, with its price descending to an approximate value of $14 USD.

The turbulent nature of cryptocurrency markets recurrently accentuates the imperative role of prudent risk management strategies. The sustainability and prosperity of decentralized finance (DeFi) applications are inextricably tied to disciplined risk mitigation efforts, and methodologies.

Appendix B

Perpetual DEX Market Manipulation

1

With every innovative feature introduced, the dual-edged sword of potential growth for the protocol and heightened exposure to potential attack vectors becomes a tangible reality. This concept is not foreign to the decentralized finance (DeFi) ecosystem. Cointelegraph estimates that such vulnerabilities led to losses exceeding $2.8 billion within the DeFi landscape during 2022 alone. Exploits typically manifest when malign entities manipulate asset prices to facilitate lucrative transactions. Let's scrutinize a tangible manifestation of such an exploit.

1.1 Case Study: Mango Markets

On the 11th of October, 2022, Mango Markets, a perpetual exchange platform operating on the Solana blockchain, fell prey to an adversarial trader. The malefactor skillfully exploited the platform's protocols, abusing the available leverage offerings and the facility to obtain loans against unrealized profits or losses (PnL). Consequently, the antagonistic agent managed to extract over $100 million in funds.

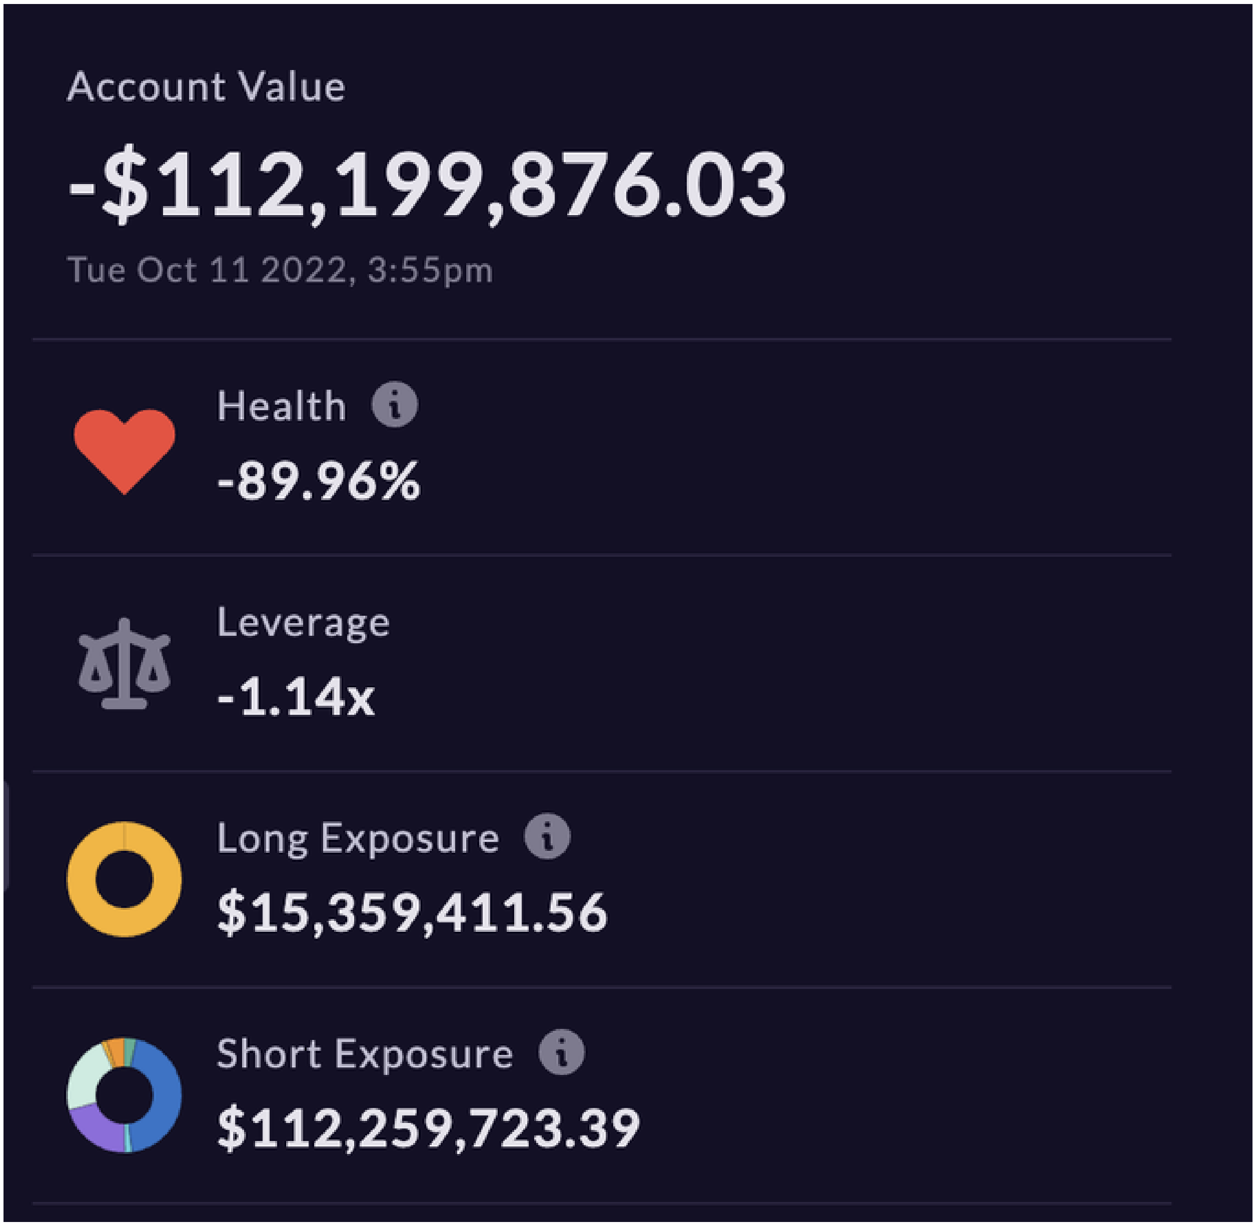

Figure B.1: Status Mango Markets perpetual exchange after the attack was completed. Source.

The successful execution of this exploit can be partly attributed to the relatively thin liquidity of the $MNGO token. In tandem with the leverage positions on the perpetual exchange, this facilitated the attacker's ability to inflate and deflate the price of MNGO, thus securing substantial profits. This incident serves as an invaluable, albeit costly, lesson underscoring the pivotal role that market parameters play in safeguarding both traders and the protocol itself.

Appendix C

About Chaos Labs

Chaos Labs is a cloud-based platform that develops risk management and economic security tools for decentralized finance (DeFi) protocols. The platform leverages sophisticated and scalable simulations to stress test protocols in adverse and turbulent market conditions. By partnering with DeFi protocols, Chaos Labs aims to create innovative solutions that enhance the efficiency of DeFi marketplaces.

The Chaos Labs team exhibits exceptional talent and represents diverse expertise, encompassing esteemed researchers, engineers, and security professionals. Chaos Labs has garnered its experience and skills from renowned organizations, including Google, Meta, Goldman Sachs, Instagram, Apple, Amazon, and Microsoft. Additionally, the team boasts members who have served in esteemed cyber-intelligence and security military units, further contributing to their unparalleled capabilities.

You can explore our past and ongoing projects for customers like Aave, GMX, Benqi, dYdX, Uniswap, Maker, and more in the Research and Blog sections of our website.

Appendix D

Acknowledgements

We express deep gratitude to Professor Dan Galai for his patient review and useful critiques of this research work.