Estimating Investor Preferences for Blockchain Security

Nir Chemaya[a], Dingyue Liu[a]

[a] *University of California, Santa Barbara, Department of Economics, 2127 North Hall, Santa Barbara, CA, 93106, USA*

Abstract

The use of decentralized exchange (DEX) platforms has been growing in the last few years. New Layer 2 (L2) blockchain alternatives provide better scalability and lower fees than the Ethereum blockchain (L1), but the security of L2 relative to L1 is unclear and difficult to identify. Using a structural model and a novel and comprehensive data set, we estimate investors' preferences for blockchain security on two main L2 networks, Polygon and Optimism. We find that traders anticipate an 0.68% (3.29%) chance of losing the transaction value when trading on Polygon (Optimism) compared to L1, and a considerable amount higher than the (0.01%-0.3%) transaction fee charged on each trade. Our work can be seen as empirical evidence of the trade-off between scalability, security, and decentralization, which is the biggest challenge of blockchain networks.

Keywords: Blockchain, Cryptocurrency, Decentralized Exchanges, Decentralized Finance, Security, Layer 2

JEL: D81, D83, G10, G19, G40

Email addresses: nir@ucsb.edu (Nir Chemaya), dingyueliu@ucsb.edu (Dingyue Liu)

1. Introduction

Liquidity pools are innovations in decentralized finance (DeFi). They allow for the exchange of crypto assets without the traditional centralized limit-order-book mechanism. Investors deposit tokenized assets into smart contracts.[1] They then exchange tokens from these pools according to the terms prescribed by a mechanism that determines the swapping price of each transaction. Uniswap is currently the largest liquidity pool protocol in DeFi, with a daily volume of roughly $2 billion, and total liquidity of $5 billion.[2]

Traditionally, most liquidity pools operate on the Ethereum blockchain, also known as the Layer 1 (L1). However, the DeFi landscape is evolving, and liquidity pool protocols like Uniswap have expanded their support to Layer 2 (L2) blockchains, such as Polygon and Optimism. These L2 blockchains address Ethereum's scalability challenges by offering lower gas fees and faster transaction processing[3], making them an appealing option for traders.

Ethereum's scalability is limited to processing 15-30 transactions per second, resulting in high gas fees. Gas fees are paid for the validation service made by the validators (miners). When only a small number of transactions can be validated in a given block, this can potentially creating blockchain congestion and generate high fees, as described by Sokolov (2021). In December 2021, the gas fees for swapping transactions in Ethereum were on average $93.30.[4] Conversely, L2 blockchains like Polygon can handle up to

[1] A smart contract is a self-executing contract with the terms of the agreement between buyer and seller being directly written into lines of code on the blockchain.

[2] Data source: https://uniswap.org/ on March 28, 2022 (only Ethereum).

[3] Speed of settlement/validation of the transactions also known as Finality.

[4] Gas fees for swapping are determined by the gas price and the number of gas units

10,000 transactions per second with minimal gas fees. Traders are highly sensitive to network fees and may postpone or abandon transactions when fees are high, a phenomenon documented in Easley et al. (2019). Similarly, the study by Cong et al. (2023) reveals that L2 scaling solutions offer substantial reductions in operating costs (gas fees), enhanced data accuracy, and promote decentralization by decreasing market concentration and fostering increased participation.

L2 blockchains offer traders an alternative blockchain network to execute transactions with improved conditions that could incentivize them to transition exclusively to these platforms. However, the degree to which traders migrate to L2 will depend significantly on their perceptions of the relative security provided by L2 and the original L1. A major concern in moving from L1 to L2 blockchain is the security of transactions and asset ownership. Assessing the actual risks of trading in L2 networks compared with L1 involves many different aspects, so this is not as straightforward to identify.

We divide the risks into three main categories. First, there are smart contract risks – there could be a bug in the code or hacking that affects the contract, or admin key access, all of which could contribute to a centralization problem. The second risk relates to the use of wrapped ether tokens when trading in L2 pools. Wrapped tokens represent blockchain native tokens [5] issued on a non-native blockchain, and the use of warped tokens thus

---

the smart contract uses to execute the transaction. In December 2021, the average gas price was 94 gwei. Data source: Etherscan data from swapping transactions on Uniswap.

[5] Native tokens are often used to pay gas fees or stake in DPoS systems. Ether (ETH) on Ethereum is an example of a native token.

includes liquidity risk. Finally, there are validation risks that depend on the particular blockchain's validation (consensus) technology. The main risk of the validation process, known as a 51% attack[6], occurs when someone or a group of people takes control of more than half of the validation authority of a blockchain network, thereby enabling them to create and manipulate transactions. To tackle Ethereum's scalability challenge, L2 solutions employ distinct validation mechanisms that expedite the validation process. However, this increased speed comes with potential risks.

It is difficult to determine how much riskier L2 blockchains are compared to L1 blockchains. One approach to estimating traders' perception of this risk is through surveys, but this method can be costly and may encounter validity issues due to the anonymous nature of users within the blockchain ecosystem. In our research, we propose an alternative approach by analyzing trading data from liquidity pools. This method captures traders' behavior and decisions, offering insights into their beliefs about risk.

Our inspiration for this approach comes from previous studies that have used market prices to reveal subjective beliefs. The core idea is that prices convey valuable information about people's perceptions, and by employing economic models, we can estimate these subjective views. For instance, past research has evaluated the value of statistical life (VSL) by comparing wages between riskier and safer jobs as discussed in Viscusi and Aldy (2003). The

[6]There have been several 51% attacks on blockchain networks. For example, there was an attack described in Garratt and van Oordt (2020), on Bitcoin Gold in May 2018. A more detailed explanation of these attacks in different blockchain validation technologies are provided in Sayeed and Marco-Gisbert (2019).

wage difference between these jobs reveals workers' beliefs about the value of their lives and the compensation they require to undertake risks. Likewise, prediction markets leverage prices to reveal subjective beliefs about the likelihood of events as seen in Wolfers and Zitzewitz (2006). This approach has also been applied to financial inquiries, such as explaining the equity premium puzzle by incorporating agents' subjective beliefs as Cecchetti et al. (2000) discuss.

By adopting this approach, we have developed a model that allows us to estimate traders' preferences for blockchain security using trading data. Our results shed light on how traders consider risk and adjust their behavior.

We use detailed data with more than five million transactions on L1 pools and more than 14 million transactions on L2 pools. The total swapping value of these transactions add up to more than $358 billion dollars. The data includes different kinds of pools with different token types (WETH/ETH, WBTC, UCSD, USDT, DAI) and L2 networks (Polygon and Optimism). We collected more than one year's worth of data with a significant variation in gas prices. Transactions largely sort in a systematic pattern; specifically, we observed larger transactions in L1 and smaller transactions in the L2 network. We wished to understand why traders still use L1 if the L2 has higher scalability and lower fees. And why did transactions sort in this way?

We first checked whether these results are due to the pool size,[7] and we found that this does not fully explain the sorting pattern. With the liquidity pools pricing mechanism, each transaction directly impacts the exchange rate

[7]Pool size refers to the amount of liquidity in the pool.

based on the size of the transaction relative to the pool's size. As the pool's liquidity increases, this effect decreases. However, a larger transaction leads to greater impact. L1 pools exist longer than L2 pools; they also have higher liquidity during this time, thus offering better exchange rates in relatively large transactions. In relatively small transactions, the price effect is low in both pools, and it is less expensive to trade in L2. We calculate the optimal monetary switching point for traders to trade on the L1 network instead of the L2 network, considering the exchange rate and gas fees. Empirical data supports the notion that traders switch to L1 earlier than predicted.

As security considerations related to L2 could significantly influence traders' behavior, we employ a structural model to capture these concerns. This model helps bridge the gap between monetary predictions about traders' transition to L2 and the empirical evidence, particularly regarding the threshold for switching. We've determined that other explanations, such as price accuracy, adoption costs, and the benefits of holding assets on the original blockchain (L1), fall short in fully accounting for the observed divergence between theoretical predictions and actual behavior.

According to our model, traders anticipate a 0.68% and 3.29% probability of incurring a transaction value loss when trading on Polygon and Optimism, respectively, compared to L1. This risk perception is considerable, especially when juxtaposed with the transaction fee range of 0.01% - 0.3% levied by Uniswap for each trade. To our best knowledge, this is the first study that quantifies traders' beliefs about these security considerations using trading data from DeFi platforms. Our methodology can be expanded to estimate trader perceptions on other DeFi or payment platforms.

The rest of the paper is organized as follows: Section 2 introduces Decentralized exchanges, L2 Implementations and Constant Product Market Maker (CPMM). The proposed model is described in Section 3. Section 4 introduces our data and provides summary statistics, while Section 5 provides results. Finally, Section 6 concludes our findings.

2. Decentralized Exchanges (DEX)

Most trading markets in the financial system are based on the traditional limit-order-book mechanism, in which buyers and sellers bid prices via a centralized organization that matches their bids. For years, cryptocurrencies and digital assets have mainly traded in centralized exchanges (CEX) such as Coinbase, which works with the limit-order book. Decentralized exchange (DEX) platforms recently entered the crypto market and have offered traders new decentralized options; the most common DEX protocols are liquidity pools.

Liquidity pools are contracts that enable agents to provide liquidity (tokens/assets) to a smart contract on the blockchain. Traders can trade tokens/assets from these pools using a pricing rule written in the code. Most of these pools use a "bonding curve" pricing rule, which is a function of the supply of tokens/assets in the pool and is also known as Constant Product Market Maker (CPMM). These pools incentivize agents to provide liquidity and become liquidity providers by giving them a swapping fee for each swapping action from the pool. These swapping fees are around 0.01% - 1%, depending on the protocol and tokens/assets of the pool. Most pools have two tokens/assets that traders can swap.

The most prominent DEX platforms are Uniswap, Sushiswap, Balancer, and Bancor. This paper will focus on the Uniswap protocol, which is the largest one available. Most of these protocols work on the Ethereum blockchain (the L1 network). Recently, some liquidity-pool protocols such as Uniswap have started to support L2 networks, such as Polygon, Optimism, Arbitrum and Celo.[8] The Uniswap protocol was initiated in November 2021 and December 2021 to support swapping on the Optimism and Polygon networks. We use Polygon and Optimism for our analysis as the alternative L2 networks for Ethereum (L1).

Recently, many researchers have explored DEX platforms in various directions. Some works (see, e.g., Lehar and Parlour, 2021; Barbon and Ranaldo, 2021) compare the various aspects of CEX and DEX platforms, such as liquidity provision, absence of arbitrage, price efficiency, and transactions cost. Additionally, many papers (e.g., Park (2021); Capponi and Jia (2021)) have introduced the CPMM mechanism and discussed its properties and conceptual flaws. We instead explore how agents decide which network to use on DEX platforms, as well as the security aspect of those decisions.

2.1. L2 Implementations and Security

Measuring the actual risks of trading in L2 networks compared to L1 involves many different aspects and so it is difficult to identify. First, to trade in DEX, traders need to use a smart contract that involves some risks of having a bug in the code, hacking into the smart contract, and admin keys access,

[8]Due to data limitations, we could not collect data from Arbitrum network. Celo is in its early stage, so we only show our analysis from Polygon and Optimism networks.

which could create a centralization problem, as mentioned in Tsankov et al. (2018); Schär (2021). Integrating Uniswap methods (codes) with different blockchain networks and token types can create different security risks.

The second risk of trading in L2 compared with L1 is the use of wrapped ether tokens when trading in L2 pools. Wrapped tokens represent blockchain native tokens issued on a non-native blockchain. While using the L2 network, traders must use the wrapped tokens of Ether (Ethereum native token) to trade this token in L2. These wrapped tokens include liquidity risk, which depends on the wrapped token-issuing mechanism (Caldarelli, 2021). The recent case of the Ronin network hack, which led to the loss of more than $600M, contributed to shedding light on these risks.[9]

Finally, the validation risks depend on the blockchain's validation (consensus) technology. To address the scalability problem of Ethereum (L1), L2 solutions use a different validation mechanism, which allows them to provide higher scalability and lower gas fee. Vitalik Buterin, one of the co-founders of Ethereum, already has identified that the biggest challenge of blockchain networks is achieving a decentralized payments system with high scalability and security. The main problem is that there is a trade-off between the three (decentralized, scalability, and security), and there is no technology that includes all the features together (known as the blockchain trilemma or scalability trilemma[10].

With that in mind, L2 implementations try to provide higher scalability and lower fees, but this has some drawbacks. There are many different L2

[9] Data Source: BBC: https://www.bbc.com/news/technology-60933174

[10] See BIS (2022) and Makarov and Schoar (2022))

solutions, each using a different approach. In our paper, we will focus on Polygon and Optimism, given our data set. Polygon is a side chain network with its native token (Matic) and validation mechanism (Proof-of-Stake), which means that the security is separate from the L1 network. Polygon is pegged to the Ethereum blockchain system, and users can transfer tokens from Polygon to Ethereum and vice versa using a bridge (see Thibault et al. (2022)).

Optimism uses a different L2 solution approach called optimistic rollups. In a rollup system like Optimism, transaction execution is moved to L2, and the data from these transactions are published on L1. Every Optimism transaction has two costs: An L2 (execution) fee and an L1 (data posting) fee. Most of the time, these fees are significantly lower than on the L1. Optimistic rollups use an "optimistic" validation approach in which the aggregators (who execute transactions on L2 and post them on L1) do not ask for proof of validity for each transaction execution. It means that the network supposes that the aggregators' transactions are valid. Another group of players, called verifiers, are monitoring the data published by the aggregators to deter any issues. A more detailed explanation of L2 implementations is provided in Thibault et al. (2022).

The bottom line is that L2 solutions use a different validation process than L1; therefore, it is difficult to tell how much riskier they are than L1. This paper aims to shed some light on how trades react to the trade-off between scalability, security, and decentralization.

2.2. Constant Product Market Maker (CPMM)

Another difference between CEX and DEX, besides being decentralized, is the pricing mechanism; in CEX, the asset price is determined by the bids of the buyers and sellers, while in most DEX platforms it is determined by the pricing formula called the constant product market maker (CPMM). The CPMM formula works so that the product of the amount of tokens X and Y in the pool must remain the same. Let's consider a liquidity pool that contains x tokens of token X and y tokens of token Y [following the notation of Barbon and Ranaldo (2021)].[11] The CPMM pricing rule means that for any time t the product of the available tokens (X and Y) in the pool equals a constant k, which can be expressed as

$$x_t y_t = k$$

The amount of both tokens in the pool at time t determines the current pool price P^t which can be expressed as

$$P^t = \frac{y^t}{x^t}$$

Let f denote the protocol swapping fee which goes for the liquidity providers and $\varphi = 1 - f$ is what left for the trader to swap. If at time

[11] This model and our extension are based on Uniswap V2 (Adams et al., 2020), in which traders can trade from the pool without any liquidity restrictions. However, Uniswap V3 (Adams et al., 2021) works in different mechanisms, providing liquidity with some price limits. Unfoutently, due to data limitations, we cannot analyze our data with the new V3 mechanism. However, Chemaya and Liu (2023) shows that the V2 model can provide highly accurate results for V3 data.

t+1 a trader wants to swap $\Delta(x)$ tokens X for getting Y tokens, we can calculate how many tokens Y $\Delta(y)$ she will get. CPMM states that

$$k = (x^t + \varphi\Delta x)(y^t - \Delta(y))$$

Solving for $\Delta(y)$:

$$\Delta(y) = y^t \frac{\varphi\Delta x}{x^t + \varphi\Delta x}$$

Further we can calculate the price of this swap,

$$P^t(\Delta x, y^t, x^t) = \frac{\Delta y}{\Delta x} = \frac{\varphi y^t}{x^t + \varphi\Delta x}$$

This formula's convexity relation implies that once traders have more demand for token X relative to token Y, the supply of this token in the pool will decrease, and thus its swapping price will increase. Additionally, this also implies that larger transactions have a larger price impact. However, the price impact would be small when the pool size is relatively large to the transaction size, as shown in Lehar and Parlour (2021).

These are essential properties of the liquidity pools that traders need to know. Once a trader can trade the same tokens X and Y in different networks, L1 or L2, the pool's size on each network could have a different price effect, one factor which will determine where the trader will choose to trade. The following section provides an extended model which allows the trader to pick which network they are willing to trade in.

3. Model

We follow the notations of Barbon and Ranaldo (2021) with an extension where trades can choose which network (i) they are willing to trade on.

Let i denote the blockchain network type, $i = 1$ is Ethereum (L1) network and $i = 2$ is Polygon or Optimisim (L2) network. Let X denote token 1, and Y denote token 2. f^i = The protocol swapping fees at the network i and $\varphi^i = 1 - f^i$ is what left for the trader to swap. T^{it} = the gas fee of swapping in network i at time t.

Given that the gas fee is paid by native tokens (Matic for Polygon and Ether for Ethereum and Optimisim[12]) for each network, we will calculate these fees in US dollars units in our data analysis to have one unit of account. By the CPMM, we can calculate how many tokens Y Δy^i trader will get when she trades on network i, which can be expressed as:

$$\Delta y^i = y^{it} \frac{\varphi^i \Delta x}{x^{it} + \varphi^i \Delta x} \tag{1}$$

Let P^{it} represents the price of making a transaction of value Δx for swapping token Y on network i.

$$P^{it}(\Delta x, x^{it}, y^{it}) = \frac{\Delta y^i}{\Delta x} = \frac{\varphi^i y^{it}}{x^{it} + \varphi^i \Delta x} \tag{2}$$

To scale our model so that we have only one unit of account for each transaction (token 1 - token 2) or (token 2 - token 1), we calculate the total value left for the trader after the swapping in token 2 units[13], which can be expressed as,

$$P^{it}(\Delta x, x^{it}, y^{it}) \cdot (\Delta x)$$

To explain how agents behave in an environment where they can choose

[12] Gas fees in the Optimisim are paid by Ether tokens. For more info: https://www.optimism.io/

[13] We will later convert them to US dollar values to have one unit of account.

which network they are willing to trade, we specify a model in which agents need to maximize their utility when choosing between swapping in the L1 network (ETH) or L2 network (POLY or Optimisim). This maximization problem should consider two main aspects: how many token Y(X) traders get from swapping token X(Y) on each network and how many gas fees they pay. On top of that, we can add a behavioral parameter of traders' beliefs about the security of each network.

Representative agent maximization problem:

$$\max_{i=0,1}\{i\cdot\pi_1\cdot u(P^{1t}(\Delta x, x^{1t}, y^{1t})\cdot\Delta x - T^{1t}) + (1-i)\cdot\pi_2\cdot u(P^{2t}(\Delta x, x^{2t}, y^{2t})\cdot\Delta x - T^{2t})\} \tag{3}$$

Where T^{it} is the gas fees in each network at time t, $P^{it}(\Delta x, x^{it}, y^{it})$ represents how many token Y(X) traders get from swapping token X(Y) on each network (a function both of the transaction size Δx and the pool size (x^{it}, y^{it}) in each network), and π_i traders' beliefs of the probability of not losing ones' transaction wealth in network i, everything is scaled to be in US dollars units.[14]

Thus, our representative agent would choose network i if and only if

$$\pi_i \cdot u(P^{it}(\Delta x, x^{it}, y^{it}) \cdot \Delta x - T^{it}) \geq \pi_j \cdot u(P^{jt}(\Delta x, x^{jt}, y^{jt}) \cdot \Delta x - T^{jt})$$

We assume our representative agent is risk-neutral[15] and maximizes the

[14]Our data resource allow us to convert everything to US dollars value and have one unit of account.

[15]There is strong evidence from many different researchers, as summarized in BIS (2022), that most of the traders in the crypto markets are risk-seeking. Assuming that the representative agent is risk neutral is a conservative assumption for our belief elicitation.

expected payoff.

$$u(v) = v$$

There are two networks, L1 (Ethereum) and L2 (Polygon or Optimism), agent choose L2 network if and only if

$$\pi_{L2}\cdot(P^{L2,t}(\Delta x, x^{L2,t}, y^{L2,t})\cdot(\Delta x) - T^{L2,t}) \geq \pi_{L1}\cdot(P^{L1,t}(\Delta x, x^{eth,t}, y^{L1,t})\cdot(\Delta x) - T^{L1,t})$$

Let $w = \Delta x$ and $P^{i,t} = P^{i,t}(\Delta x, x^{i,t}, y^{i,t})$ we can further write:

$$\begin{aligned}
\pi_{L2} \cdot (w \cdot P^{L2,t} - T^{L2,t}) &\geq \pi_{L1} \cdot (w \cdot P^{L1,t} - T^{L1,t}) \\
(\pi_{L2} \cdot P^{L2,t} - \pi_{L1} \cdot P^{L1,t})w &\geq \pi_{L2} \cdot T^{L2,t} - \pi_{L1} \cdot T^{L1,t} \\
w &\leq \frac{\pi_{L2} \cdot T^{L2,t} - \pi_{L1} \cdot T^{L1,t}}{\pi_{L2} \cdot P^{L2,t} - \pi_{L1} \cdot P^{L1,t}} \\
w &\leq \frac{\pi_{L1} \cdot T^{L1,t} - \pi_{L2} \cdot T^{L2,t}}{\pi_{L1} \cdot P^{L1,t} - \pi_{L2} \cdot P^{L2,t}}
\end{aligned}$$

$$w = \frac{\pi_{L1} \cdot T^{L1,t} - \pi_{L2} \cdot T^{L2,t}}{\pi_{L1} \cdot P^{L1,t} - \pi_{L2} \cdot P^{L2,t}} \tag{4}$$

This is the representative agent's threshold transaction size in which she will switch from the L2 network to L1.

Consider when π_i in each network are equal, meaning there's no security concerns of L2 relative to L1, $w^\star$ represents the optimal threshold at which the representative agent should switch from trading in the L2 network to Ethereum. At any given time t with given pools sizes and gas fees, we can calculate the theoretical $w^\star$, which we will discuss this in more detail in section 5.1.

$$w^* = \frac{T^{L1,t} - T^{L2,t}}{P^{L1,t} - P^{L2,t}} \tag{5}$$

When the representative agent's empirical threshold, $\hat{w}$, is smaller than $w^{\star}$, it means agents are switching to Ethereum even though it is less profitable. Section 5.3 will provide a comparison between $\hat{w}$ and $w^{\star}$ and a robustness test to check if $\hat{w}$ is statistically significantly smaller than $w^{\star}$. This deviation will be captured by the security parameter in our model. We can estimate the representative agent's beliefs on security, the chance of losing the transaction value when trading on L2 compared to on L1, as,

$$S_{L2,L1} = 1 - \frac{\hat{\pi_{L2}}}{\hat{\pi_{L1}}} = 1 - \frac{P^{2t}(\hat{w}, x^{2t}, y^{2t}) \cdot \hat{w} - T^{2t}}{P^{1t}(\hat{w}, x^{1t}, y^{1t}) \cdot \hat{w} - T^{1t}} \tag{6}$$

this will be discussed in section 5.3.

4. Data

This paper collects transactions from a total of 21 liquidity pools on the Uniswap V3 protocol,[16] including 8 in the L1 network (ETH) and 13 in L2 network (6 in POLY, 7 in OPT). We follow this selection process: first, we collected trading data from all networks and pools, and we identified some key features of the pools, like average transactions per day, tokens involved, pool transaction fee, etc. Then we sorted pools that were only available on both L1 and L2, shared the same liquidity pool fee, and shared the same token types. Finally, we chose from only pools with more than 200 average daily transactions, and with a minimum of 70 or more transactions per day.

These pools allow traders to trade the same token types in L1 and L2

[16]The Uniswap protocol is a peer-to-peer system designed for exchanging cryptocurrencies (ERC-20 Tokens).https://docs.uniswap.org/protocol/introduction

(POLY and OPT) and have a sufficient amount of transactions per day. These pools jointly contribute 63% of the transactions on the Uniswap when considering pools that are available for trades to trade the same pair of tokens on L1 and L2.[17] Six different tokens are swapped in these pools (DAI, USDC, USDT, WMATIC/MATIC, WBTC, WETH/ETH), and the pool fee ranges from 0.01% to 0.3%.

Our main analysis will focus on three liquidity pools, with one from the three networks (ETH, POLY, OPT). Each of the three pools have the same pair of tokens (USDC and WETH/ETH),[18] and have the same protocol swapping fees $f^1 = f^2 = 0.05\%$. Those pools are the biggest ones in our data set and contribute more than 50% of the daily transactions. Utilizing blockchain explorer services (Uniswap Data Extractooor),[19] we are able to track each and every Erc-20 tokens transactions that happened in the liquidity pools.[20] We collected data from December 22, 2021, the launch date of the Polygon network pool, until December 31, 2022.[21]

Within this 12 months period, we obtained a total of $2,789,976$ swapping

[17]Many pairs of tokens are network specific and can be traded only on one of the networks. For example, on 12/05/2022, only 53.36% of the transactions on Ethereum pools were with tokens that were available on L2 (Data source Uniswap Data Extractooor); there was a similar situation with L2 pools: only 55.39% (46.56%) of the transactions on POLY(OPT) pools were with tokens that were available on L1.

[18]The Ether (ETH) tokens on the L2 blockchains are wrapped tokens (WETH).

[19]https://www.uniswap.shippooor.xyz/

[20]The ERC-20 introduces a standard for Fungible Tokens, in other words, they have a property that makes each token be exactly the same (in type and value) as another token. https://ethereum.org/en/developers/docs/standards/tokens/erc-20/

[21]Optimism pool lunching day was one month before on 11/12/2021.

transactions (exchange between USDC and WETH/ETH) from L1 ETH network, 4, 991, 764 swapping transactions from L2 POLY network, and 4, 323, 672 swapping transactions from L2 OPT network.[22] That is a total of more than 12 million swapping transactions, which resulted in a sum of $237 billion. The distributions of the amount swapped in the three platforms differ during the time of interest (Figure 1). We modified the magnitude of the values for large numbers, so this graph is more readable. All transactions with a value greater than 10,000 dollars are over-written to 10,000 (for this graph only). We observe that the L2 distribution is right-skewed for most smaller transactions (on POLY less than $923, on OPT less than $64). On the other hand, the majority of the transactions on L1 ETH are larger than $5, 002, five to a hundred times more than on L2.

We also calculated the daily mean gas fee for swapping in each network.[23] During this time period, the average daily mean gas fee for swapping was $22.93 on L1, $0.559 on OPT, and $0.030 on POLY.

We also collected data on the daily size of the three liquidity pools from Uniswap. ETH's pool had a higher pool size during our observation dates, with an average of $277.6M, while the average size of POLY was $13.40M, and the average size of OPT was only $4.72M. Figure 3 is the time series presentation of the three liquidity pools' size during this time period.

[22]ETH Contract address: 0x88e6a0c2ddd26feeb64f039a2c41296fcb3f5640.

POLY Contract address: 0x45dda9cb7c25131df268515131f647d726f50608.

OPT Contract address: 0x85149247691df622eaf1a8bd0cafd40bc45154a9.

[23]Data Source: Etherscan, Polyscan & Optimistic.etherscan.

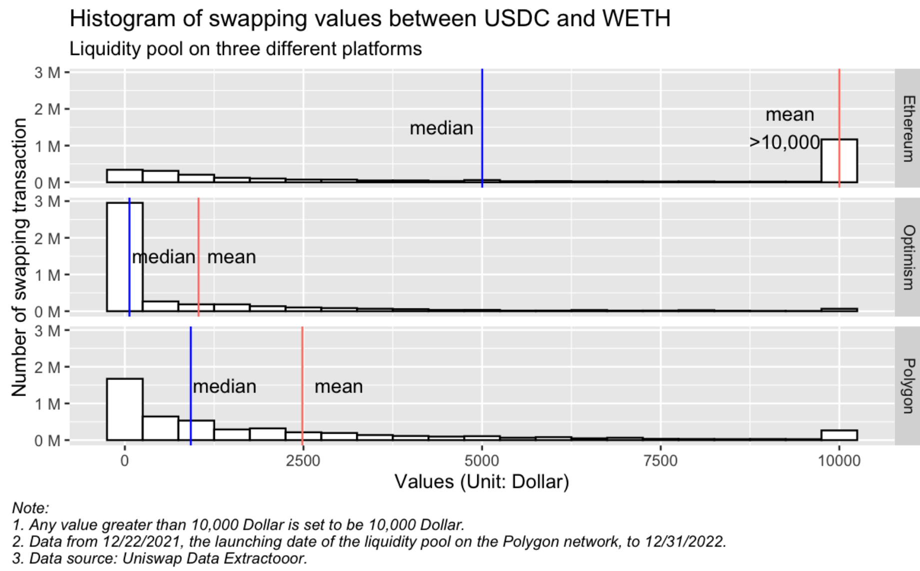

Figure 1: Histogram of swapping values between USDC and WETH on three platforms

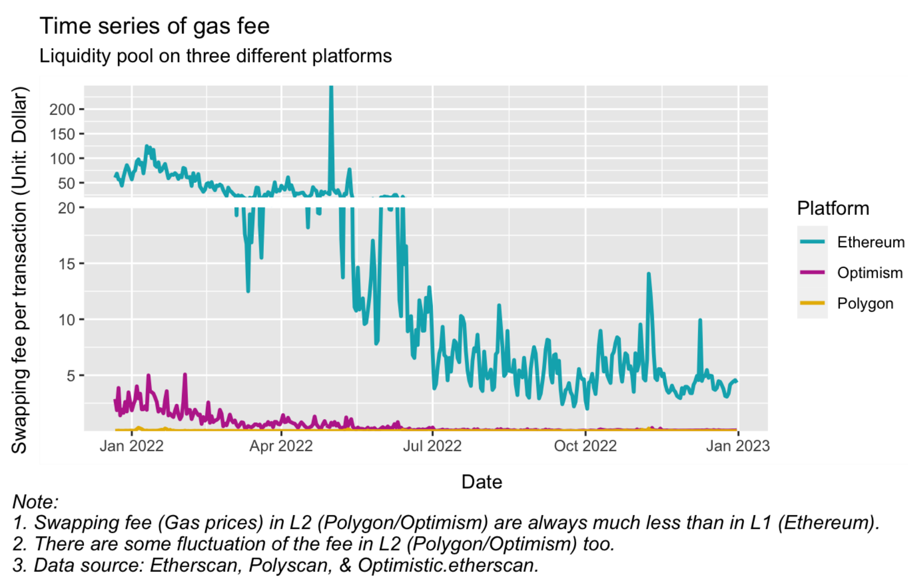

Figure 2: Time series of gas fee

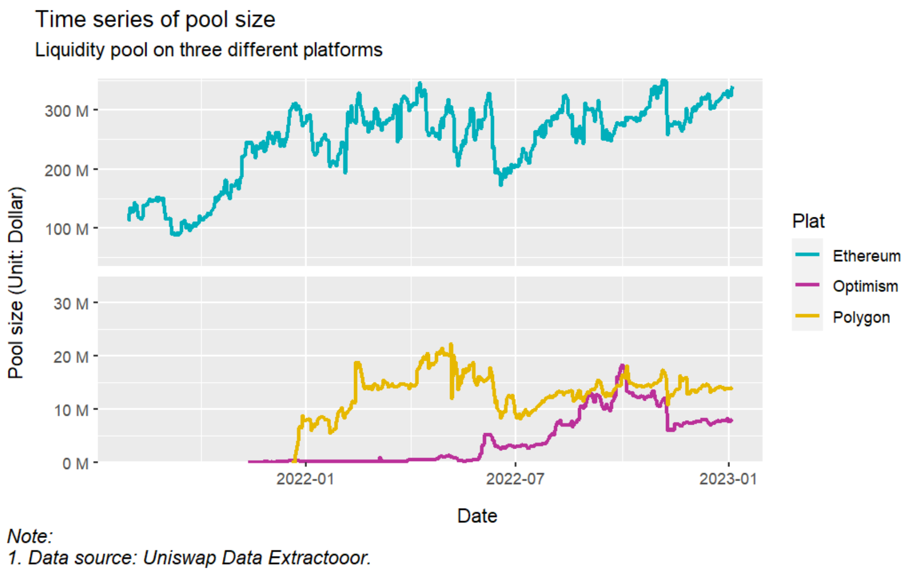

Figure 3: Time series of pool size

5. Results

This massive trading data from liquidity pools capture traders' behavior and decisions. In this section, we will use our model to analyze that data and estimate traders' preferences for blockchain security. We will mainly discuss the result from the three pools introduced in section 4, which are the biggest ones in our data set. We will first show details of estimating the security of the POLY (L2) network relative to ETH (L1) using POLY & ETH, USDC/WETH 0.05% pools; then we will show results from OPT (L2) network and other pools.

We estimate traders' belief in security by studying behavior deviation from monetarily optimal decisions. In section 5.1, we calculate the monetarily optimal switching point W^*, as a function of liquidity pool size and transaction fee. Section 5.2 presents our empirical strategy for finding the actual switching point $\hat{W}$ from the data. Section 5.3 summarizes the security concern from W^* and $\hat{W}$ based on our model. Finally, section 5.4 further shows our model's generalizability to other networks and pairs of tokens in different pools.

5.1. Monetarily optimal switching point W^*

Why do traders still use L1 if the L2 has higher scalability and lower fees? The exchange rate of a swapping transaction in the liquidity pool is determined by the liquidity pool size and the size of the transaction itself (recall Equation 2 in the model section). The higher liquidity in L1 makes L1 pools have a lower price impact, offering better exchange rates for relatively big transactions. On the other hand, if the transaction size is relatively small, the price effect in both pools is low; considering L1 has a higher gas fee, it would then be cheaper to trade in L2. Given pool size and gas fee at a low swapping size, it would be cheaper to swap on L2; at some switching point, it would be optimal to switch to swap on L1.

Take April 11th, an arbitrary day, as an example: the ETH, USDC/WETH 0.05% pool size is \$322.93M, and the POLY, USDC/WETH 0.05% pool size is \$18.58M. The mean gas fee is \$27.37 per swapping transaction on ETH, and \$0.02 per swapping transaction on POLY. Following Equation 2, while also taking into consideration the gas fee, we calculate the total value left for the trader after the swapping. As shown in Figure 4,[24] it is better to swap on Polygon at first, and then better on Ethereum once the swapping value becomes larger (to be exact, once the swapping value is larger than $\$16,442$).

[24]Using R package *ggforce* (Pedersen, 2021).

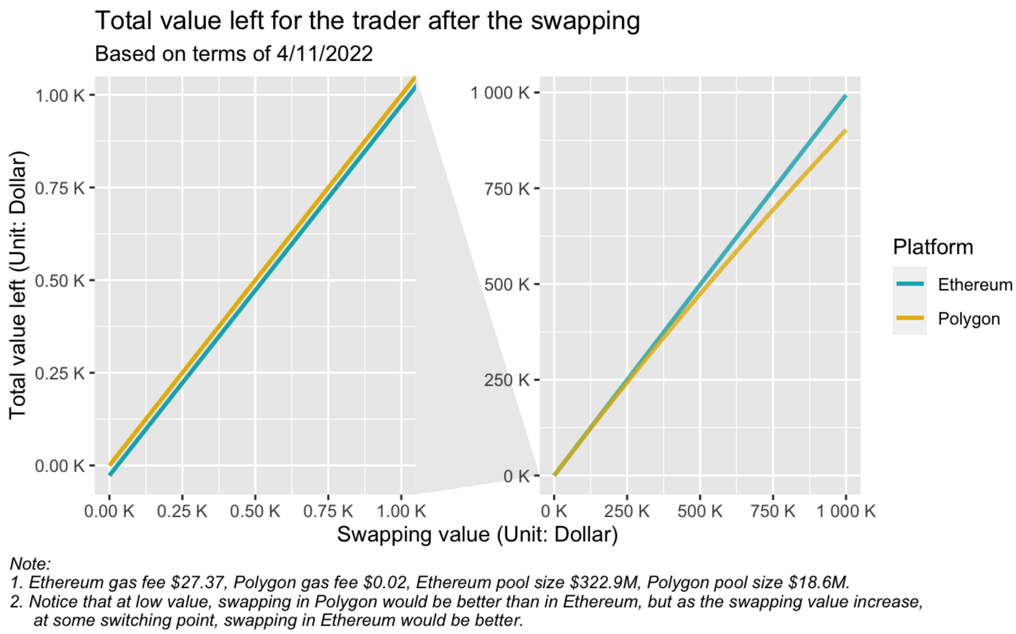

Figure 4: Total value left for the trader after swapping

We calculate this monetarily optimal switching point W^* for all dates in our data, and get Figure 5.

5.2. Empirical switching point $\hat{W}$

The monetarily optimal switching point W^* can explain some reasoning behind traders trading on both platforms and separated in a certain way, yet empirical data supports that traders are switching to L1 for much lower transactions.

In order to find the representative agent's empirical threshold from the data, given that we are facing a binary classification problem, we implement a binary logistic model.[25]

[25] A competing method is linear discriminant analysis, a linear method in classification,

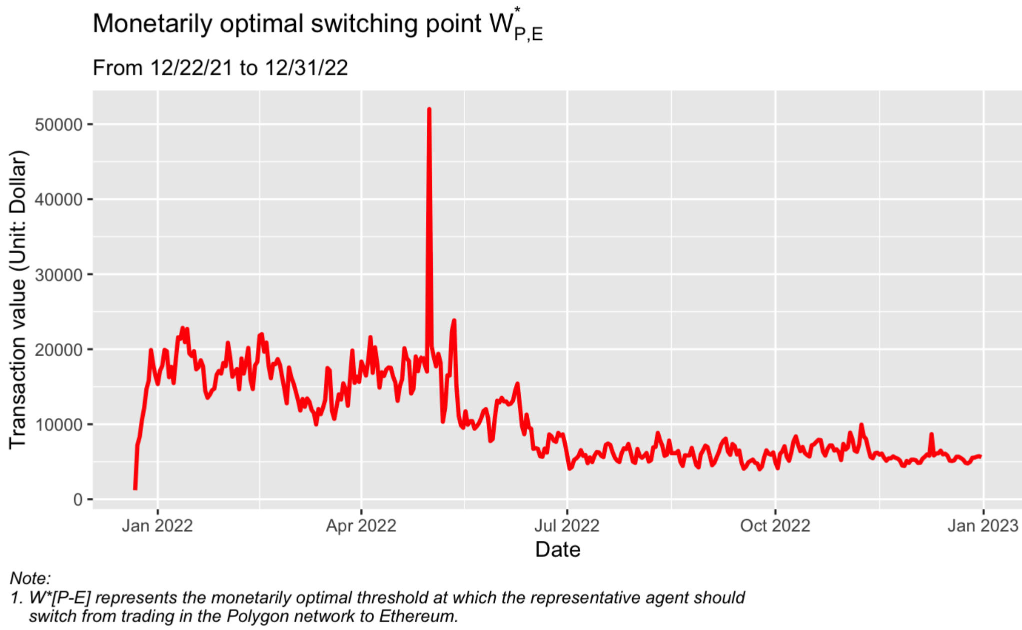

Figure 5: Monetarily optimal switching point $W^*_{P,E}$

A binary logistic model states that the probability of outcome Y belongs to class y given predictor W equal to a logistic function.

$$Pr(Y = 1|W) = \frac{e^{\beta_0 + \beta_1 W}}{1 + e^{\beta_0 + \beta_1 W}}$$

There are two classes, transaction is on L2 ($Y = 1$), and transaction is on Ethereum ($Y = 0$). Our predictor W is the transaction value (in Dollar unit). It is a linear model as the logit, or log-odds, is linear in W.

$$log(\frac{Pr(Y = 1|W)}{1 - Pr(Y = 1|W)}) = \beta_0 + \beta_1 W$$

We report the the summary of the logit regression result using 2022/4/11

see Appendix A for more discussion.

data in Table 1.[26] The coefficient for W (Transition value) is negative, meaning the higher the transaction value is, the lower the log odd, that is, the lower the probability of this transaction is on L2 (and higher probability is on Ethereum). Both the interception and the coefficient for W are statistically significant.

	Logit Regression
(Intercept)	1.72^{***}
	(0.02)
Transaction value	$-1.278e^{-4***}$
(in Dollar)	$(2.822e^{-6})$
AIC	18401.68
BIC	18417.62
Log Likelihood	-9198.84
Deviance	18397.68
Num. obs.	21364

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 1: Binary Logistic Model result: 2022/4/11 data

To find the empirical threshold value $\hat{W}$, we obtain $\hat{\beta}_0$ and $\hat{\beta}_1$ from the regression, and the best threshold probability $\hat{Pr}(Y = 1|W)$.[27]

$$\hat{W} = \frac{1}{\hat{\beta}_1}(log(\frac{\hat{Pr}(Y = 1|W)}{1 - \hat{Pr}(Y = 1|W)}) - \hat{\beta}_0)$$

[26]Using R (R Core Team, 2020), and R package *texreg* (Leifeld, 2013).

[27]See Appendix B for more details.

Figure 6 shows a time series of the calculated $\hat{W}$. On 2022/4/11, this empirical threshold is \$3,469.

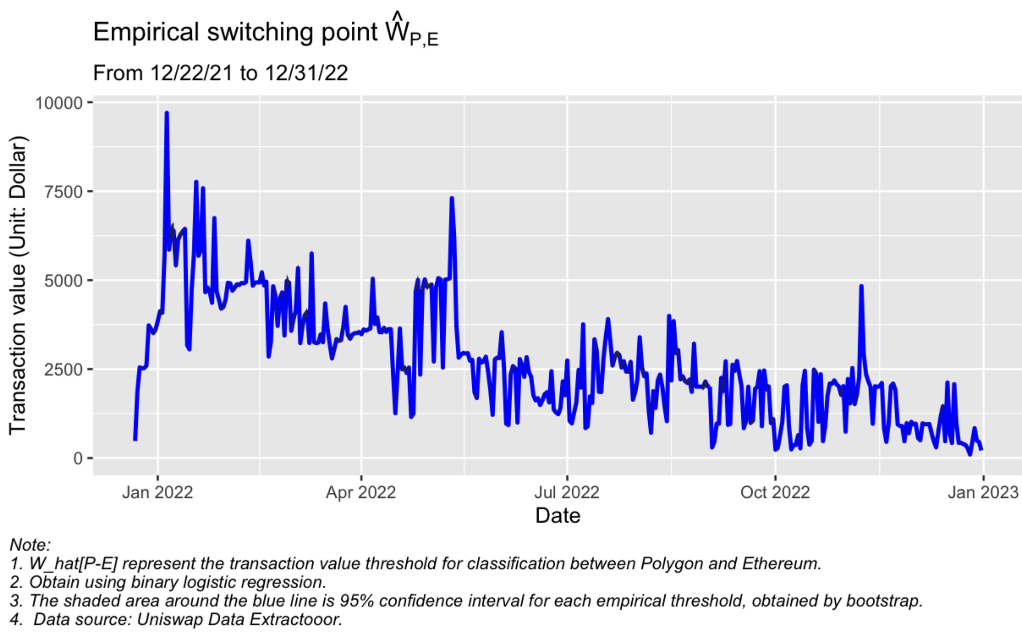

Figure 6: Empirical switching point $\hat{W}_{P-E}$

5.3. Estimating belief on security

Our structure model captures traders' security concerns about L2. This security concern can explain the gap between the switching point from the pure monetary prediction W^* and empirical $\hat{W}$. Figure 7 is a direct comparison of the two time series in one graph. The shaded area around the blue line represent the area of 95% confidence interval obtained by running bootstrap on the transaction data.[28]

[28] Using R package *boot* (Davison and Hinkley, 1997).

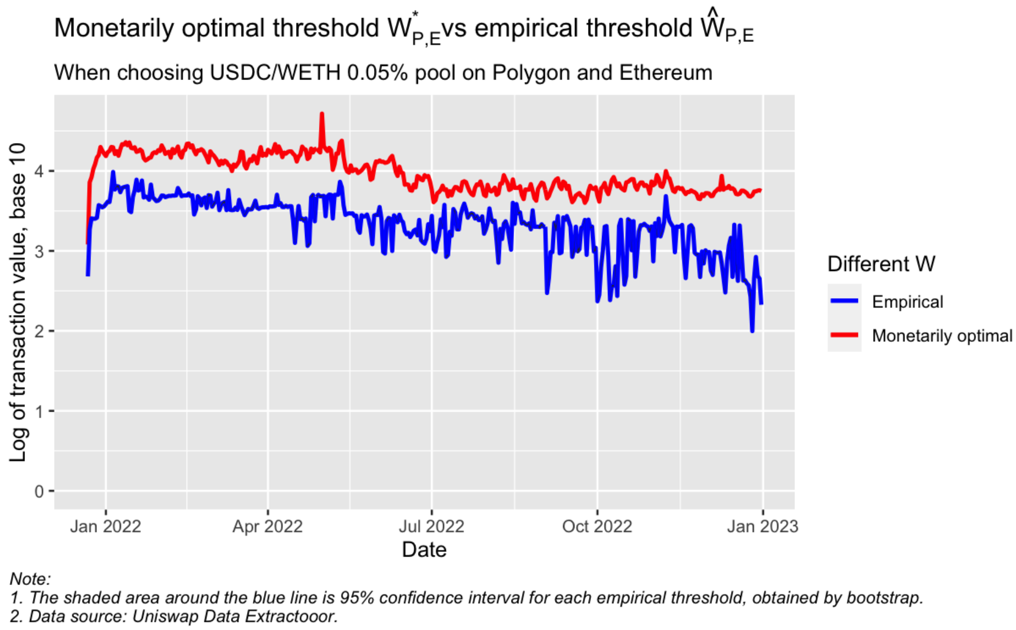

Figure 7: Monetarily Optimal Threshold $W^*_{P,E}$ vs Empirical Threshold $\hat{W}_{P,E}$

Notice that the monetarily optimal switching point W^* is always above the empirical threshold $\hat{W}$. This is consistent with our prediction of the model. The intuition is as follows, due to the security concerns, people would switch from trading in Polygon network to Ethereum network earlier, thus the gap we observe from empirical evidence and model prediction, that is, the gap of $\hat{W}$ and W^*.

As defined in section 3, $S_{L2,L1}$, our estimator of the representative agent's beliefs on security, is the chance of losing the transaction value when trading on L2 compared to on L1. $S_{L2,L1}$ should be greater than 0, as the probability of not losing ones' transaction wealth in L1 network should always be smaller than on L2 network, since Layer-2 network building on Layer-1 network. The calculation confirms this as the estimator is always greater than 0.

Here in Figure 8, the y axis is $S_{L2,L1} = S_{P,E}$; the higher the estimator, the more security concerns our representative agent holds on the Polygon network. The mean of the analysis time period is 0.751%, suggesting that in this period, on average, agents think there is 0.751% more chance of losing transactions on Polygon compared to Ethereum. The median is 0.554%.

The ratio is significantly different from 0 (greater than 0), as its 95% confidence interval, obtained by running bootstrap on the transaction data, never cover 0.

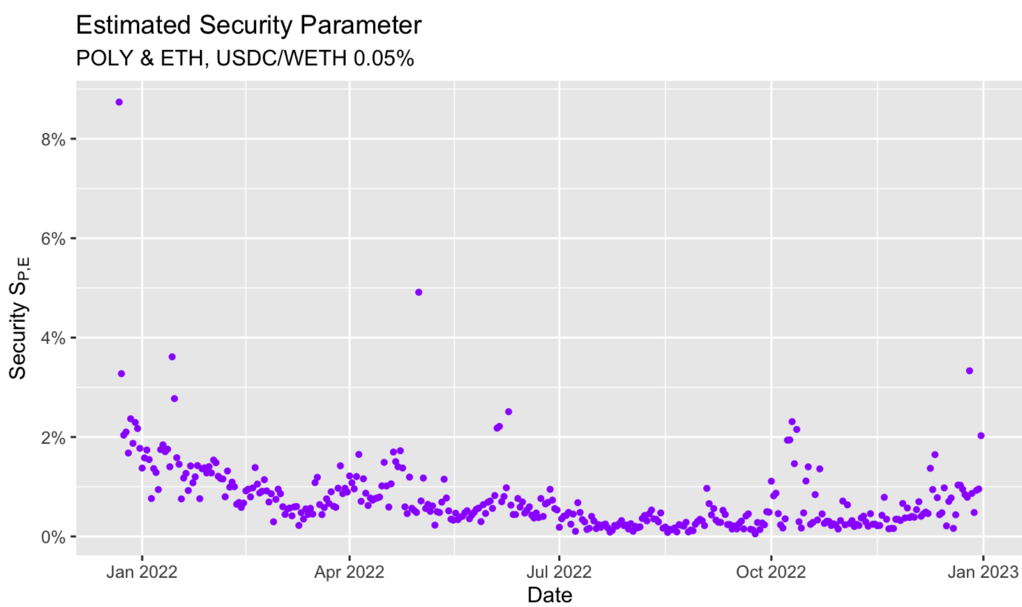

Figure 8: Estimated Security Parameter

5.4. Investigating Alternative Factors to Explain the Results

In this section, we will explore alternative explanations that could account for the observed gap between the monetary optimal switching point and the empirical one. Specifically, we will closely examine potential explanations related to factors such as price accuracy, adoption cost, and the advantages of owning assets on L1. However, we find that these explanations are less likely to account for the observed gap, as the available data does not provide strong support for them.

Moreover, taking into account the blockchain trilemma, as discussed in 2.1, which highlights the trade-off between scalability and security, users who transition to more scalable networks like L2 should be mindful of this trade-off. In light of this, we assert that security assumes a pivotal role in explaining the observed gap and our research findings.

5.4.1. Price Accuracy

One argument suggests that the gap between the monetary optimal switching point and the empirical one may be attributed to different prices across networks. However, our data indicates that prices across networks, particularly in large pools with well-known tokens, are equal between L1 and L2. Figure 9 plots the time series of WETH price relative to USDC in the 0.05% fee liquidity pool for all three platforms. Prices across L1 ETH, L2 POLYGON, and L2 OPTIMISM are almost identical such that the time series overlaps.

This suggests that arbitrageurs operate across layers, potentially holding L2 and L1 accounts and periodically transferring funds between networks to avoid incurring cross-chain costs. This concept resembles arbitrageurs trading on both centralized exchanges (CEX) and decentralized exchanges (DEX), where moving funds from DeFi to CeFi can be costly. The presence of arbitrageurs across networks weakens the argument that traders do

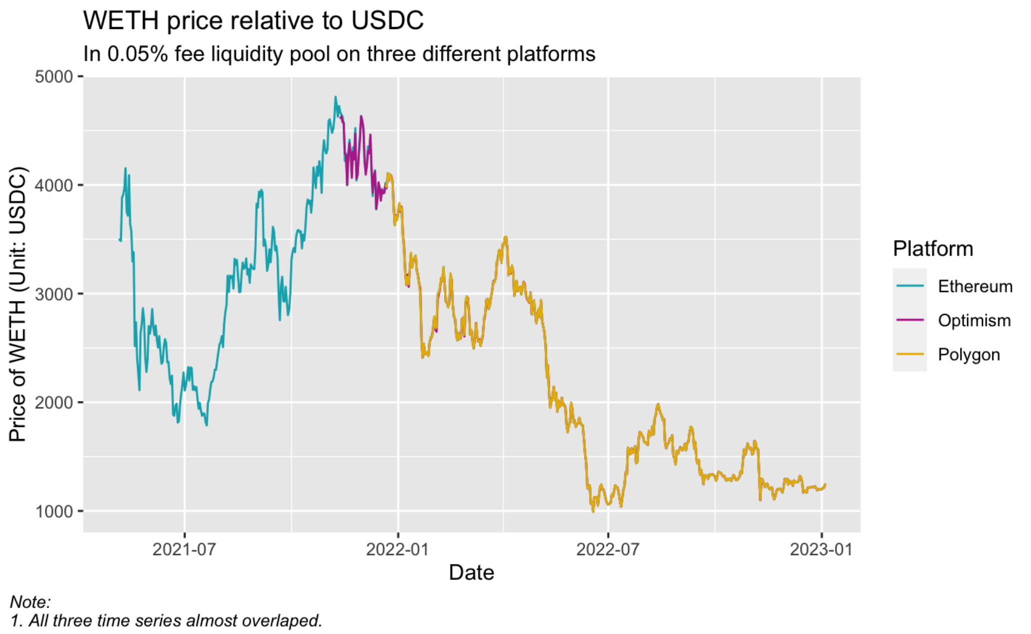

Figure 9: WETH/USDC price in 0.05% fee liquidity pool

not monitor prices across networks, and the potential gap could be due to monitoring costs.

However, if this were the case, we would anticipate greater price volatility between networks, and the security parameter might exhibit more noise or even negative values. In light of our model, when prices on L1 are more favorable than L2, price inaccuracy should manifest as a lower security parameter and potentially even as negative values, signifying that L2 is more secure than L1. Nevertheless, our data consistently demonstrates positive security parameters, indicating the accuracy and consistency of exchange prices on L2 in comparison to L1.

5.4.2. Adoption Cost

The use of L2 solutions entails an adoption cost initially, which may influence users' transition from L1 to L2. To adopt L2 solutions, users need to transfer funds from L1 to L2 (involving bridging mechanisms) and create a new wallet on L2, requiring familiarity with the L2 network. Although L2 solutions aim to streamline the process by enabling the use of the same digital wallet and wallet ID across L1 and L2 networks, the adoption cost may still be substantial for certain users, depending on users' level of sophistication and familiarity with these systems.

Given that our model employs a representative agent framework, it might overlook this adoption cost concern, and it is possible that small users find the adoption cost relatively affordable while wealthier users face higher barriers. To assess this argument, we conduct analyses on a subset of the data, namely, on wallets that traded on both L1 and L2 during the period of analysis. This allow us to examine whether users who hold funds in both L1 and L2 exhibit the same pattern of smaller transactions on L2 and relatively larger ones on L1, as well as whether their switching point is lower than the optimal one.

Table 2 gives an overview of the subset data. 17,246 unique wallet addresses swapped on both ETH and POLY for the USDC/WETH 0.05% pools. Together they contribute to about 5%-7% of the total transactions we observed in this period, and about 1%-6% of the total transaction value on each platform.

We then conduct the same analysis as in 5.3 to obtain the mean (median) estimated security parameter. Our findings not only support that individual users follow the same patterns as our representative agent, but also highlight

	DEX	# TXN (%)	Total Volume (%)	Mean Estimated Security Parameter (Median)
Full	ETH	2,789,976	$220,065,992,854	
	POLY	4,991,764	$12,401,731,851	0.751% (0.554%)
Subset	ETH	142,538 (5.1%)	$2,119,468,116 (1.0%)	
	POLY	346,088 (6.9%)	$ 697,737,558 (5.6%)	1.672% (1.250%)

Table 2: Subset Data for users that swap on both ETH and POLY

that the representative agent result is a lower bound of this security estimate, thereby weakening the adoption cost argument. A similar plot to Figure8 can be found in Appendix C for this subset data, FigureC.17.

5.4.3. Benefit of Owning Assets on L1 vs. L2

In addition to the liquidity risk associated with holding tokens on L2, which is one of the security concerns we highlight, our model assumes that tokens on L2 and L1 are essentially the same. However, it is plausible that tokens possess different utility values. Users may utilize their tokens on L1 in other DeFi applications that generate higher returns compared to L2, thereby providing additional value to L1 tokens and potentially explaining the observed gap.

Nevertheless, it is worth noting that numerous DeFi applications are currently available on L2, including DEX platforms and lending protocols, offering a diverse range of financial options that can occasionally yield higher returns than holding tokens on L1. For example, our data reveals that providing liquidity on L2 generates higher daily returns compared to L1 (Figure 10). Despite this, liquidity providers still exhibit a preference for providing liquidity on L1. This behavior reinforces our security concern, as it suggests that the switching cost from L1 to L2 may be relatively low compared to the higher returns from providing liquidity on L2, yet liquidity providers still choose L1 as their preferred option.

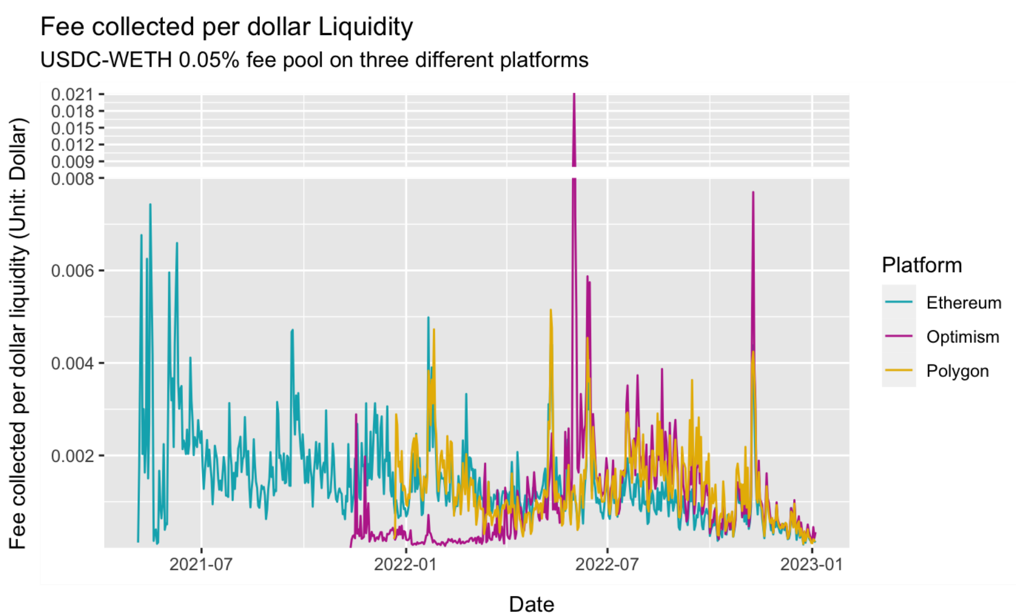

Figure 10: Fee collected per dollar liquidity for USDC-WETH 0.05% liquidity pool

5.5. Generalization

The analysis conducted in the above sections can be generalized to other L2 network, and also other liquidity pools with different tokens.

We first look at the corresponding pool in Optimism mentioned in section 4. Appendix C documented the corresponding Figure 7 and Figure 8 for this same pair of tokens and same protocol swapping fees, but comparing Optimism and Ethereum. Similar to our results from swapping behavior in the Polygon pool case, again, there's always a gap between the empirical threshold and monetarily optimal threshold in Figure D.18. The estimated security parameter for Optimism of this token pair (Figure D.19) has more fluctuation, and on average, higher security concerns from the representative agent compared to Polygon (mean is 3.53% and median is 2.64%). We will discuss this difference between Polygon and Optimism later.

To further verify our model, we performed the same exercise to 5 other token pairs on Polygon and 6 other token pairs on Optimism. These pools have different tokens involved, and also have various protocol swapping fees. Table 3 summarizes all the liquidity pools we have analyzed.

Some of these pools were not available until recently; in order for us to compare our model result across different pools and networks, we take the intersection of period of time for all pools, which is August 5, 2022 to December 31, 2022, a period of 149 days. Appendix D provides the histogram of swapping values and the estimated security parameter for these additional 11 pools. All previous analysis is replicated, as we observe the histogram of swapping values sort in the same way as we previously introduced, with small transactions in L2 and large transactions in L1. Moreover, the gap between

Pool fee	L2 Polygon	L2 Optimism
DAI/USDC	0.01	0.01
DAI/WETH		0.05
MATIC/WETH	0.3	
USDC/USDT	0.01	0.01
USDC/WETH	0.05	0.05
USDC/WETH	0.3	0.3
USDT/WETH		0.05
WBTC/WETH	0.05	0.05

Table 3: List of liquidity pool analyzed

the empirical threshold and monetarily optimal threshold are all robust, and the estimated security parameters are also similar (within the same network).

Based on each day's trading volume in each pool, we assign a weight to calculate the weighted mean of the security parameter $S_{L2,L1}$ for the two L2 platforms. The weighted mean for Polygon is 0.68%, and for Optimism is 3.29%, meaning that the representative agent believed that there is 0.68% more chance of losing transactions on Polygon compared to Ethereum, while there is 3.29% more risk on Optimism compared to Ethereum.

As one might notice from Appendix D, there's more fluctuation and worse security estimator for Optimism than Polygon; indeed, as shown in Figure 11, the Optimism's weighted mean is significantly greater from Polygon's at 95% level. This suggests traders believe that the Optimism network is less secure than Polygon. Perhaps the "optimistic" approach in the validation process of optimistic rollups reduces the reliability from a trader's point of view.

Another possible explanation is that the Community of the Polygon network is much greater than the Optimism, which generates more reliability of the traders in the Community and less about the technology behind it. But this is speculation and a topic for future research.

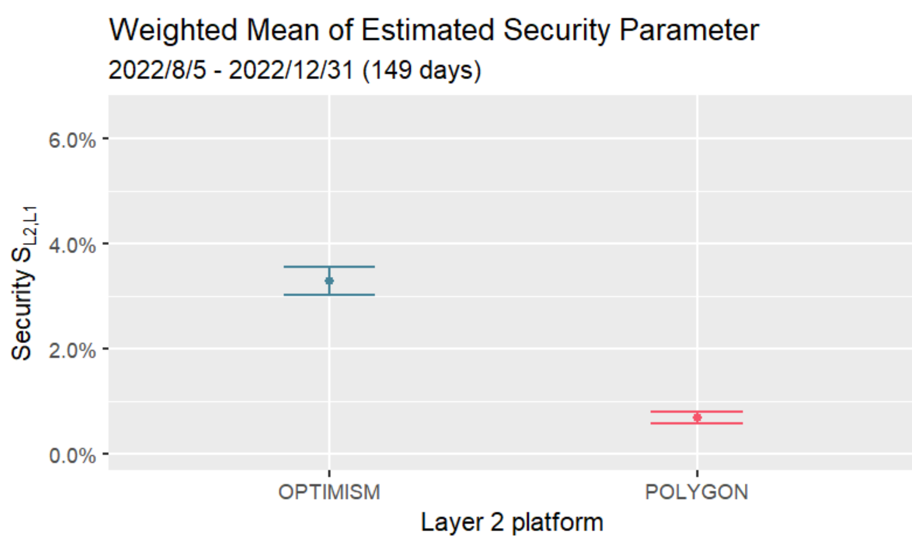

Figure 11: Weighed Mean of Estimated Security Parameter

6. Summary

The primary focus of our analysis of this novel trading environment in DEX platforms is to quantify traders' belief about security issues regarding L2 compared to L1. To do so, we analyzed trading data using a structural model. Our model calculates the monetarily-optimal switching point for traders to trade on the L1 network instead of L2. Empirical data supports the idea that traders use L1 for lower transactions even though it is less expensive to trade on L2. We argue that security concerns have a critical role in explaining this gap.

Our model reveals that, on average, traders anticipate a 0.68% (3.29%) chance of losing transaction value when trading on Polygon (Optimism) compared to L1, which is a substantial risk considering the (0.01%-0.3%) transaction fees charged per trade. Our analysis utilized a large and diverse dataset that incorporated various gas prices, different types of tokens, and two L2 networks (Optimism and Polygon). Despite this variation, we consistently obtained similar results, which highlights the robustness of our findings. Moreover, we have rigorously established that alternative explanations such as price accuracy, adoption costs, and the advantages of holding assets on L1 are less influential in explaining the observed preference for L1 over L2.We also develop preliminary insights on the impact of L2 solutions to the financial inclusion of DEX. L2 solutions, allow traders with low stakes to enter a market with a low-gas-fee environment. The number of swapping transactions on L2 is much higher than on L1, and these are mostly small-size transactions. The high gas fees in L1 do not allow traders with a small budget to trade when gas fees are high relative to their small transactions. Our work can be seen as empirical evidence of the trade-off between scalability, security, and decentralization, which is the biggest challenge of blockchain networks.

Looking forward, our novel methodology can be applied to other L2 networks, allowing researchers to estimate traders' security concerns across different networks in DEX platforms. The long-term effect of the introduction of the L2 networks is yet to be explored. Will concerns about security be reduced when L2 has been in existence longer? DeFi markets should be explored further, and we invite more researchers to study this new environment.

7. Acknowledgment

The authors would like to thank Rod Garratt, Erik Eyster, Gary Charness, Peter Rupert, Kieran Walsh, Sevgi Yuksel, Ryan Oprea, Guo Yu, Anthony Zhang, Bobby Pakzad-Hurson, Jennifer Kao, Giulio Caldarelli, Danny Klinenberg, Toshio Ferrazares, and participants at seminars at UC Santa Barbara, UC San Diego, UC Irvine for helpful comments and suggestions. The authors thank Sio U Wong, Yutong Chen, Jonas Chen, Eric Liu, Elia Ching, Alok Thakrar, Cynthia Wang, and Sean Lee for their excellent research assistance in collecting the data for this research.

8. Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix A. Linear Discriminant Analysis

A competing method of binary logistic regression is linear discriminant analysis, a linear method in classification. While the relative efficiency of linear discriminant analysis (LDA) is superior to binary logistic regression (BLR) if the LDA's assumptions are met (Efron, 1975), the assumption of normality is hard to meet with our data. In one predictor (W) case, the LDA assumes that $W|Y = k \sim N(\mu_k, \sigma^2)$, that is, the predictor given a different class, follows a normal distribution with different mean and variance. We test this assumption and found the predictor is far from normal distributed through Skewness-kurtosis graph (Cullen et al., 1999).[29]

Appendix B. Threshold the predicted probability

After obtaining the logit regression, we can predict the class (transaction is on L2 or Ethereum) by thresholding the predicted probability. For example, one might predict $Y = 1$ (on L2) for any transaction value whose predicted probability is greater than 0.5. Or, if we are being conservative in predicting transaction value to be in Ethereum, we could predict $Y = 1$ (on L2) for any transaction value whose predicted probability is greater than 0.1.

To evaluate the classification performance under different threshold probability, one can construct confusion matrix and pin down the threshold probability that obtain a low false positive rate (FPR, the fraction of negative examples that are classified as positive, which in our study is the portion of transaction on Ethereum that are classified as on L2) while also maintain-

[29]Using R package *fitdistrplus* (Venables and Ripley, 2002).

ing a low false negative rate (FNR, the portion of transaction on L2 that are classified as on Ethereum). We want to choose the probability threshold that is closest to (FPR, FNR) = $(0, 0)$. There are many ways to determine which threshold probability corresponds to the smallest distance, but we calculate the euclidean distance between each point of (FPR, FNR) and $(0, 0)$. Figure Figure B.15 and Figure B.16 show the ROC curve and optimal threshold selection for 2022/4/11 data.[30] The optimal threshold for that day is 77.91%, meaning when predicting platform, only when the $Pr(Y = 1|W) > 77.91\%$, we category the transaction to be on L2.

[30]Using R package *ROCR* (Sing et al., 2005).

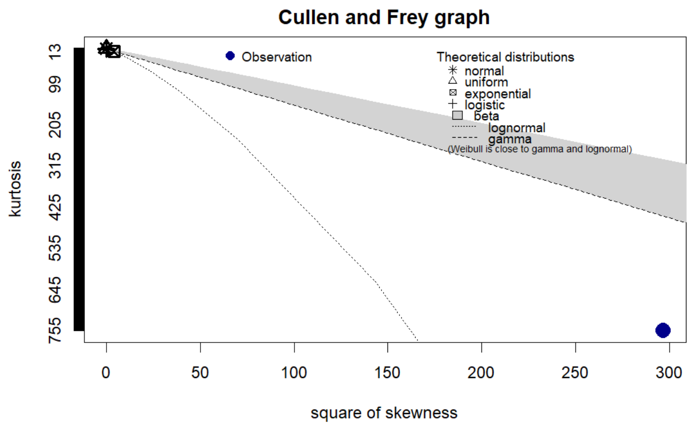

Figure A.12: Skewness-kurtosis graph, for Ethereum transactions

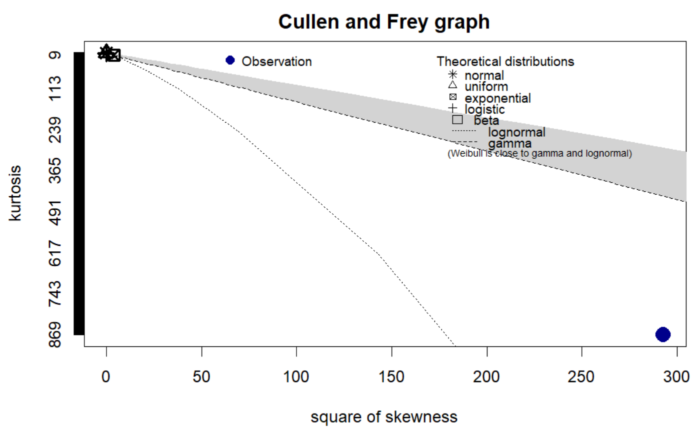

Figure A.13: Skewness-kurtosis graph, for Optimism transactions

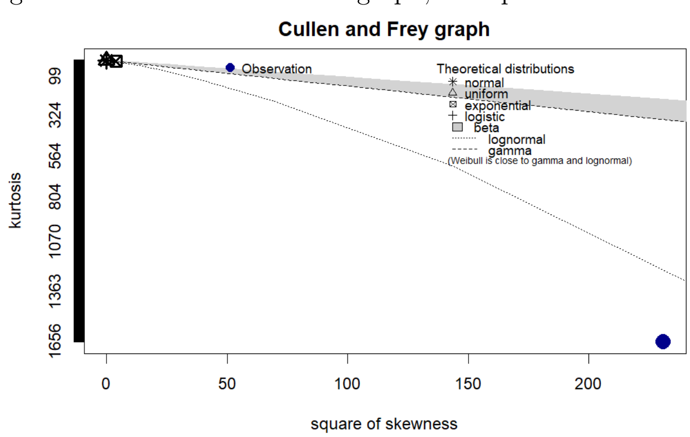

Figure A.14: Skewness-kurtosis graph, for Polygon transactions

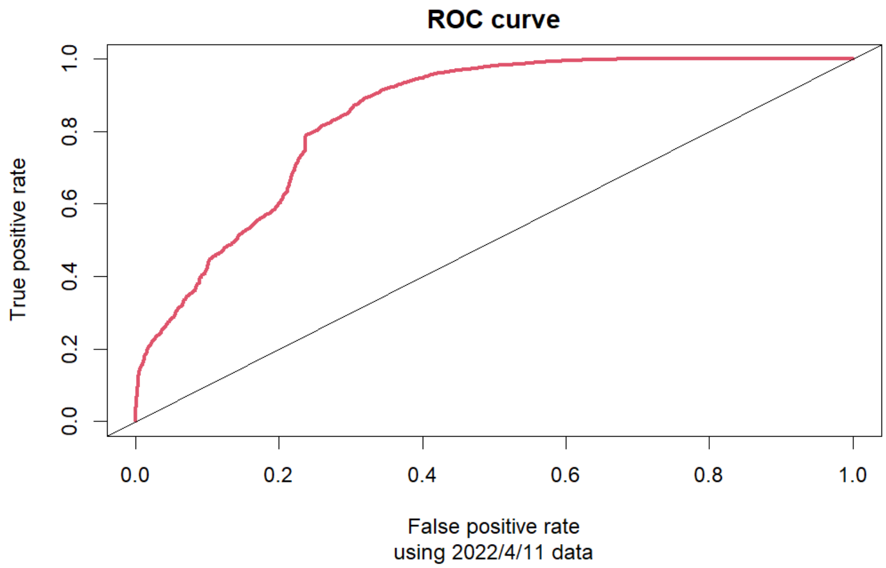

Figure B.15: Example ROC curve (2022/4/11 data)

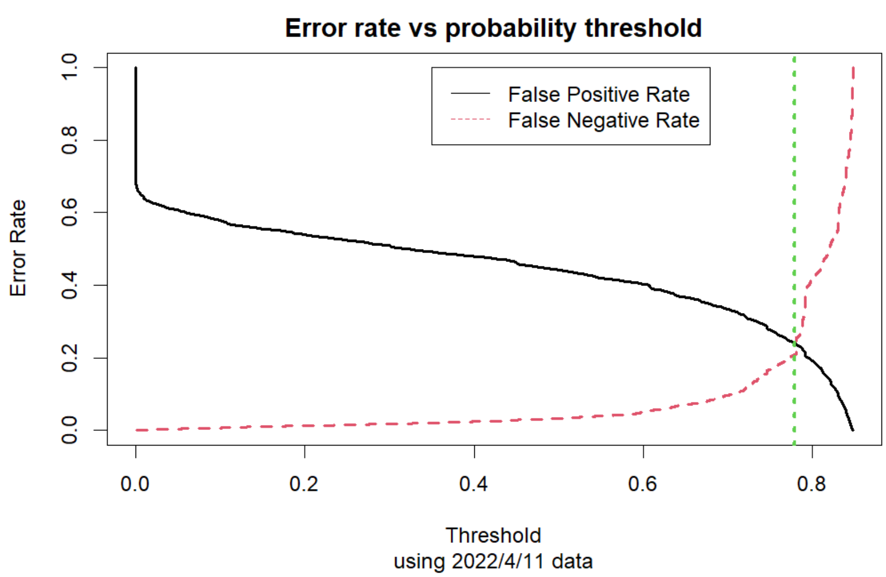

Figure B.16: Example FPR, FNR graph (2022/4/11 data)

Appendix C. Wallets that swap on both ETH and POLY

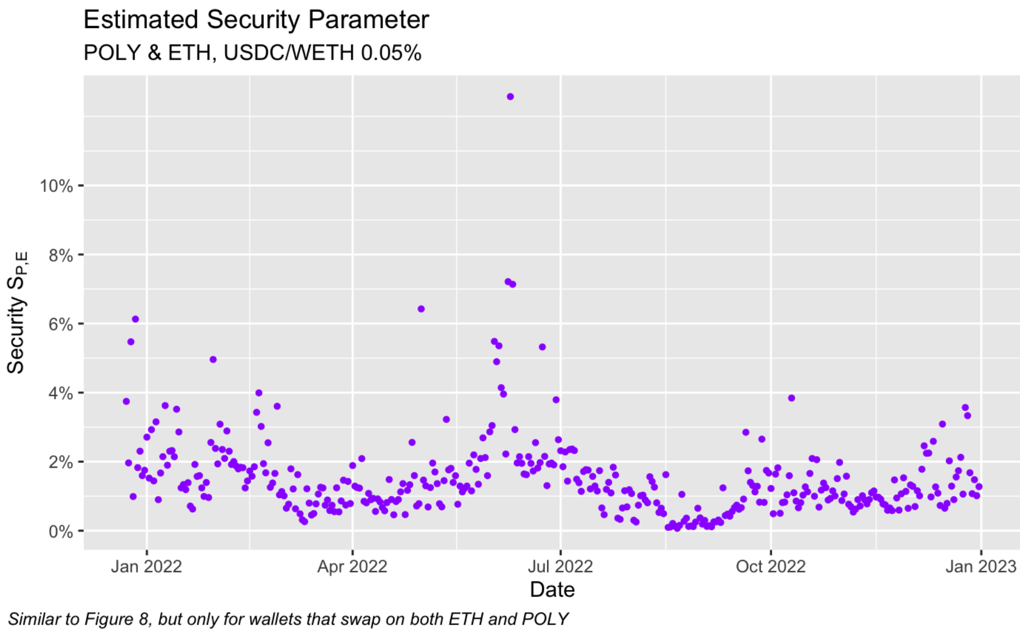

Figure C.17: Estimated Security Parameter for subset data

Appendix D. Generalization Optimism

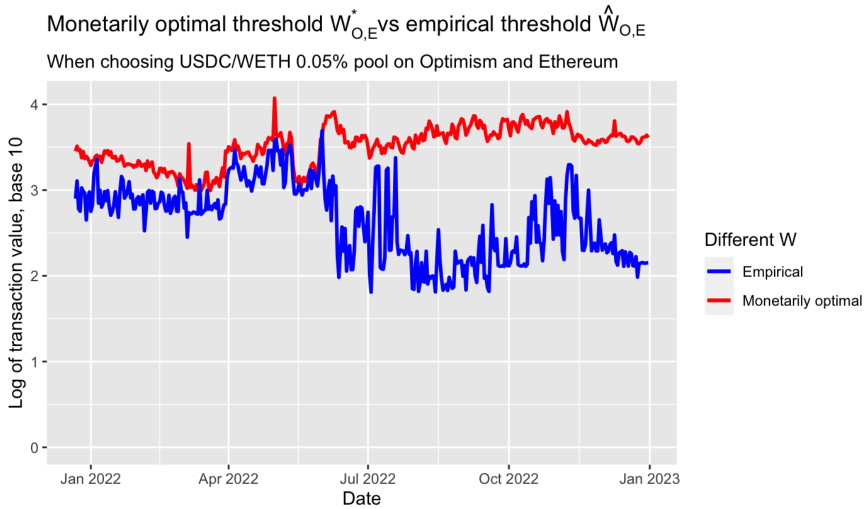

Figure D.18: Monetarily Optimal Threshold $W^*_{O,E}$ vs Empirical Threshold $\hat{W}_{O-E}$

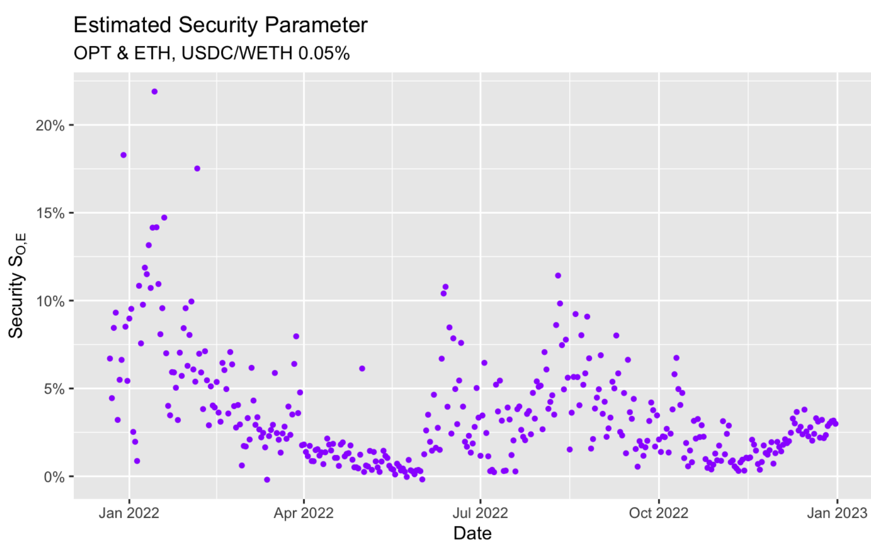

Figure D.19: Estimated Security Parameter, O,E

Appendix E. Generalization all other pairs

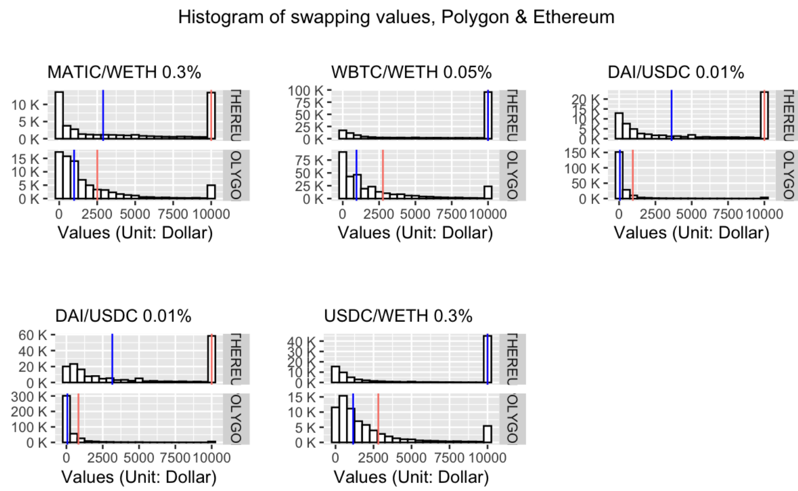

Figure E.20: Histogram of 5 other pairs from Polygon and Ethereum

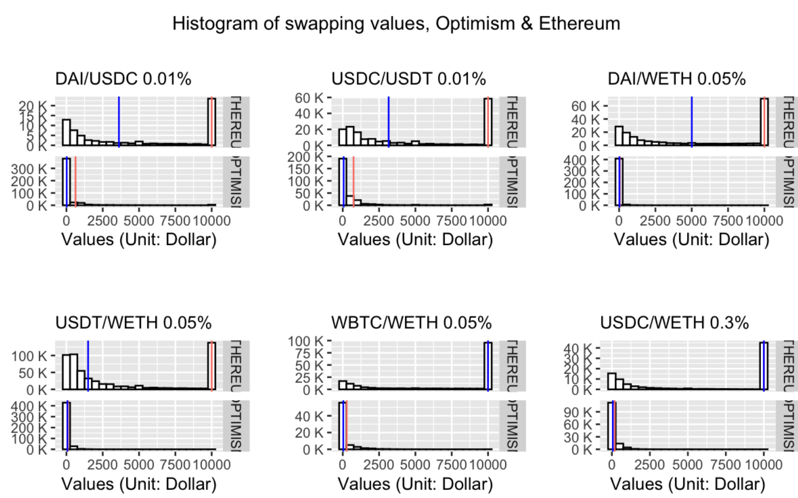

Figure E.21: Histogram of 6 other pairs from Optimism and Ethereum

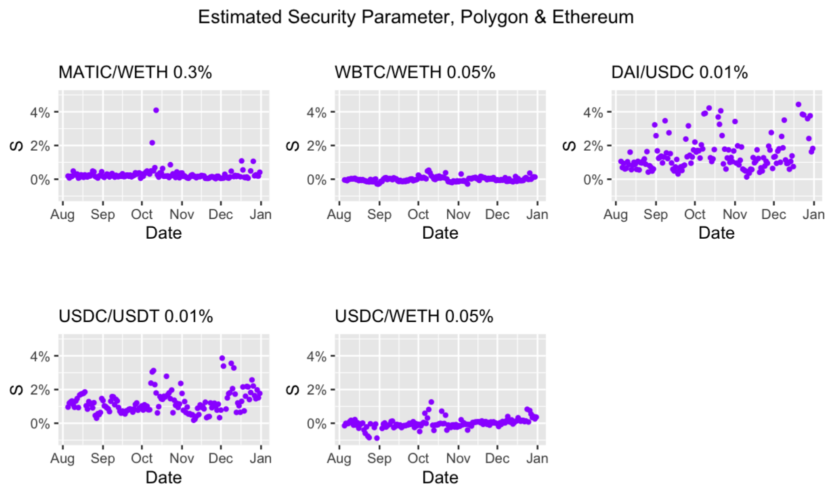

Figure E.22: Estimated Security parameter S for 5 other pairs from Polygon and Ethereum

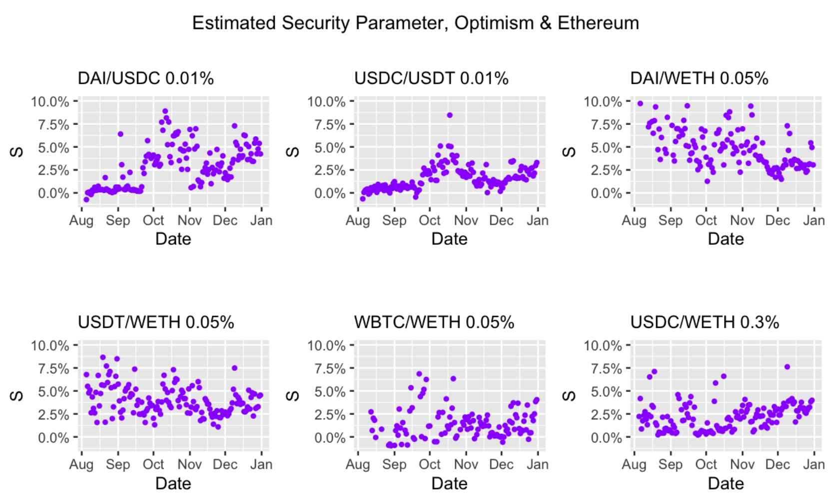

Figure E.23: Estimated Security parameter S for 6 other pairs from Optimism and Ethereum

References

Adams, H., Zinsmeister, N., Robinson, D., 2020. Uniswap v2 core .

Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D., 2021. Uniswap v3 Core. Technical Report. Technical report.

Barbon, A., Ranaldo, A., 2021. On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. arXiv preprint arXiv:2112.07386 .

BIS, 2022. "the future monetary system". Annual Economic Report 2022 .

Caldarelli, G., 2021. Wrapping trust for interoperability: A preliminary study of wrapped tokens. Information 13, 6.

Capponi, A., Jia, R., 2021. The adoption of blockchain-based decentralized exchanges. arXiv preprint arXiv:2103.08842 .

Cecchetti, S.G., Lam, P.s., Mark, N.C., 2000. Asset pricing with distorted beliefs: are equity returns too good to be true? American Economic Review 90, 787–805.

Chemaya, N., Liu, D., 2023. Analyze uniswap v3 data using v2 methods. Available at SSRN 4372372 .

Cong, L.W., Hui, X., Tucker, C., Zhou, L., 2023. Scaling Smart Contracts via Layer-2 Technologies: Some Empirical Evidence. Technical Report. National Bureau of Economic Research.

Cullen, A.C., Frey, H.C., Frey, C.H., 1999. Probabilistic techniques in exposure assessment: a handbook for dealing with variability and uncertainty in models and inputs. Springer Science & Business Media.

Davison, A.C., Hinkley, D.V., 1997. Bootstrap Methods and Their Applications. Cambridge University Press, Cambridge. URL: http://statwww.epfl.ch/davison/BMA/. iSBN 0-521-57391-2.

Easley, D., O'Hara, M., Basu, S., 2019. From mining to markets: The evolution of bitcoin transaction fees. Journal of Financial Economics 134, 91–109.

Efron, B., 1975. The efficiency of logistic regression compared to normal discriminant analysis. Journal of the American Statistical Association 70, 892–898.

Garratt, R., van Oordt, M.R., 2020. Why fixed costs matter for proof-of-work based cryptocurrencies. Available at SSRN 3572400 .

Lehar, A., Parlour, C.A., 2021. Decentralized exchanges. Technical Report. working paper, University of Calgary and University of California, Berkeley.

Leifeld, P., 2013. texreg: Conversion of statistical model output in R to LaTeX and HTML tables. Journal of Statistical Software 55, 1–24. URL: http://dx.doi.org/10.18637/jss.v055.i08.

Makarov, I., Schoar, A., 2022. Cryptocurrencies and Decentralized Finance (DeFi). Technical Report. National Bureau of Economic Research.

Park, A., 2021. The conceptual flaws of constant product automated market making. Available at SSRN 3805750 .

Pedersen, T.L., 2021. ggforce: Accelerating 'ggplot2'. URL: https://CRAN.R-project.org/package=ggforce. r package version 0.3.3.

R Core Team, 2020. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria. URL: https://www.R-project.org/.

Sayeed, S., Marco-Gisbert, H., 2019. Assessing blockchain consensus and security mechanisms against the 51% attack. Applied Sciences 9, 1788.

Schär, F., 2021. Decentralized finance: On blockchain-and smart contract-based financial markets. FRB of St. Louis Review .

Sing, T., Sander, O., Beerenwinkel, N., Lengauer, T., 2005. Rocr: visualizing classifier performance in r. Bioinformatics 21, 7881.

Sokolov, K., 2021. Ransomware activity and blockchain congestion. Journal of Financial Economics 141, 771–782.

Thibault, L.T., Sarry, T., Hafid, A.S., 2022. Blockchain scaling using rollups: A comprehensive survey. IEEE Access .

Tsankov, P., Dan, A., Drachsler-Cohen, D., Gervais, A., Buenzli, F., Vechev, M., 2018. Securify: Practical security analysis of smart contracts, in: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 67–82.

Venables, W., Ripley, B., 2002. Random and mixed effects, in: Modern applied statistics with S. Springer, pp. 271–300.

Viscusi, W.K., Aldy, J.E., 2003. The value of a statistical life: a critical review of market estimates throughout the world. Journal of risk and uncertainty 27, 5–76.

Wolfers, J., Zitzewitz, E., 2006. Interpreting prediction market prices as probabilities.