Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Improved Liquidity for Prediction Markets

Practical Project

Lukas Kapp-Schwoerer

lukaska@ethz.ch

Distributed Computing Group
Computer Engineering and Networks Laboratory
ETH Zürich

Supervisors:
Robin Fritsch, Quentin Kniep
Prof. Dr. Roger Wattenhofer

March 27, 2023

Acknowledgements

I would like to thank my supervisors Robin Fritsch and Quentin Kniep for their support and guidance throughout this project. I would also like to thank the Distributed Computing Group lead by Prof. Dr. Roger Wattenhofer for providing me with the opportunity to work on this exciting topic.

Abstract

This work discusses the rise of decentralised prediction markets facilitated by smart contracts and automated market makers (AMMs). These markets allow individuals to bet on the outcome of a diverse range of events, including politics, economics, pandemics, and many more. The paper analyses the state of prediction markets through both theoretical and data analysis, identifying shortcomings in current AMMs and proposing a novel market maker that aims to provide improved liquidity for converging markets. The theoretical section reviews mechanisms of order books, AMMs, and market resolutions, as well as game theoretic properties of prediction markets. In the data analyses we scrape and study over 2 million transactions, analysing biases and accuracy against several dimensions. Through this data analysis we identify problems with liquidity provisioning for converging prediction markets, and we propose the Smooth Liquidity Market Maker (SLMM) to address this issue. The SLMM is expected to improve liquidity provisioning, and thereby increasing both trading volume and the accuracy of prediction markets.

Contents

Chapter 1

Introduction

Humans have always been interested in predicting future events. In the Olympic games of ancient times people would bet on the outcome of athletic events. In modern times traders bet in stock markets on the future performance of companies. However, such markets were limited in their specific domains such as sports and finance because they were dependent on a centralised intermediary to organise the markets. In recent years, many decentralised applications based on Ethereum and other blockchains have emerged. One particular application class are prediction markets. These markets facilitate the prediction of future events through financial bets. The recent surge of prediction markets has enabled forecasts on a much larger variety of events such as politics, economics, pandemics, natural disasters, sports, pop culture events and military events. Many of these prediction markets use smart contracts to facilitate predictions, with the prices being determined by automated market makers (AMMs). The vision of these prediction markets is that they can aggregate the forecasts of many people and provide a more accurate prediction than any individual could alone.

In this project, we analyse the state of prediction markets through theoretical and data analysis. After finding shortcomings in current automated market makers, we finally propose a novel market maker that aims to provide improved liquidity for converging markets.

To begin with, we describe the theoretical mechanisms of order books, automated market makers and market resolutions, as well as some game theoretic properties of prediction markets in chapter 2. In a second step we scrape and analyse over 2 million transactions from the prediction markets on Polymarket[1] ans provide an analysis thereof in chapter 3. We analyse biases in this historic data, as well as compare the markets' accuracy against several other dimensions. We also perform a specialised analysis on the liquidity provisioning in these markets in section 3.5. After observing problems with liquidity provisioning for converging prediction markets in our scraped data, as well as discussing the theoretical incentive structure that can lead to this problem, we finally propose a new market making algorithm in section 4.1. This "Smooth Liquidity Market Maker" (SLMM) aims to provide improved liquidity for converging markets, thereby increasing both trading volume and the accuracy of prediction markets.

[1] `https://polymarket.com/`

Chapter 2

Theoretical Background

2.1 Definition of Prediction Markets

Prediction markets are markets in which participants make financial bets on the outcome of future events. For example, a market may include 'yes'/'no' shares, where 'yes'/'no' refer to a question such as "Will event X happen before date Y?". The owners of shares with the correct outcome receive a predetermined amount, e.g. 1 USD per share. We will now describe mechanisms for such markets in more detail.

Centralised vs. Decentralised Prediction Markets

Markets can be centralised, decentralised, or a hybrid thereof. In a centralised market, a single cooperation coordinates market creation, trades, resolution and payouts. Most literature on prediction markets was conducted at times when centralised markets were predominant, such as the Iowa Electronic Market, run by the university of Iowa as an academic experiment [1]. Since the advent of Ethereum [2] and other smart contract platforms, prediction markets with various degrees of decentralisation have been created. For example, prediction markets on the Polymarket platform[1] have smart contract based trading, liquidity provision, market making, market resolution and payouts – only the market creation remains centralised in the case of Polymarket.

Examples of Prediction Markets

Prediction Markets for some narrow domains have been around for a long time. For example, the history of sports betting is believed to date back to the 23rd Olympic Games in the 7th century BC [3]. In recent decades, some prediction markets have been created for research purposes, e.g. the Iowa Electronic Markets which was created in 1998 by the university of Iowa and allows betting on elections

[1] https://polymarket.com/

in the US [4]. There are also prediction markets at corporations, for example at Google and Ford [5, 6]. In these internal markets employees bet on events relevant for the companies. Most recently, platforms offering predictions on a wide range of domains have been created, for example Polymarket. The data analysis in chapter 3 will focus on data from Polymarket.

2.2 Mechanisms of Prediction Markets

As seen in section 2.1, a variety of prediction markets exist. In this section we focus on common mechanisms of decentralised / semi-decentralised prediction markets. In particular we will focus on mechanics used by Polymarket, since the data analysis in chapter 3 will focus on data from Polymarket.

2.2.1 Order Book Based Market Making

Markets can suffer from the so-called "thin market problem" [7]: Traders must coordinate when and on which assets to trade. This can be overcome by using an intermediary. Today, two common types of intermediaries for market making exist: *order book based market making* and *automated market makers*. In the case with an order book, market participants specify a price at which they are willing to buy or sell a specific outcome token [8]. These positions are then entered into an order-book. Another market participant may then buy or sell the outcome token at the price specified in the order book. This approach as the downside that the think market problem is only overcome to the extent that there are enough market participants willing to place orders in the order book.

2.2.2 Automated Market Makers

An alternative to order book based market making is the use of Automated Market Maker (AMM)s. An AMM is an intermediary with which a market participant may trade at any given time. The basic idea of AMMs is that they hold a pool of outcome tokens called *liquidity pool*, which they use to trade with market participants. The liquidity pool is provided by *liquidity providers*, who get fees on trades in exchange for providing liquidity. We will now describe the mechanisms of AMMs as proposed by Hanson et. al [7] and Angeris et. al [9]. This theoretical foundation is similar to the one used by Polymarket.

Buying & Selling Outcome Tokens and AMM Invariants

When a market participant wants to buy or sell a specific outcome token from the AMM, the AMM determines the price based on an *invariant*. Such an invariant is a formula that takes the number of outcome tokens held by the AMM's liquidity pool as input and outputs a number that ought to be kept constant between buy / sell operations by market makers. There are various types of variants, such as the logarithmic market scoring rule described by Hanson et. al. [7] and the more recent constant product invariant [9] giving rise to a Constant-Product Market Maker (CPMM). With CPMMs, the invariant specifies to keep the product of number of outcome tokens held by the AMM per outcome token type constant. In detail, the steps for purchasing a specific outcome token from the automated market maker are:

1. The market participant gives funds (e.g. USDC) to the AMM.
2. The AMM converts the funds into outcome tokens, with equal amounts of outcome tokens being created per possible outcome. The new outcome tokens are added to the liquidity pool of the AMM, which leads to a violation of the invariant.
3. The AMM gives the market participant outcome tokens of the requested type until the invariant is restored. Out of the funds given to the market participants some fees are deducted for the liquidity providers. The fraction of the funds provided in step 1 and the amount of tokens given to the market participant in step 3 implies the price at which the market participant bought the outcome tokens.

Similarly, when selling tokens to the AMM, the AMM receives outcome tokens of a specific type and then converts outcome tokens back into collateral until the invariant is restored. The collateral recovered during this process minus fees for liquidity providers is the remuneration for the market participant.

Liquidity Provision

AMMs have a *liquidity pool* that holds tokens of the different possible outcomes. This liquidity pool is used when a market participant wants to buy or sell a specific outcome token. Initially, the AMM is provided with *liquidity* by *liquidity providers*. This liquidity is collateral from which the AMM creates an equal amount of each possible outcome token. At a later stage in time, the AMM may have different amounts of tokens for different outcomes, due to market participants buying and selling outcome tokens. If a liquidity provider wants to add liquidity to the AMM and the AMM has an imbalance of outcome tokens, the liquidity provider does not add all newly created outcome tokens to the liquidity pool, but refunds such an amount of token to the liquidity provider that the proportions of tokens in the liquidity pool are the same as before the liquidity provider added liquidity.

Liquidity providers are incentivised to provide liquidity by earning fees on trades. For example, a market might have a fixed fee of 0.5%, such that 0.5% of all trade amounts are distributed to the liquidity providers. The distribution among the liquidity providers is determined by the fraction of so called *liquidity shares* which liquidity providers receive when adding liquidity.

On the other hand, liquidity providers are also exposed to a risk: As described earlier, buying and selling tokens upholds an invariant such as the constant product invariant. In consequence, the proportions of outcome tokens in the liquidity pool may change. For example, if the possible outcomes are 'yes' and 'no', and market participants buy outcome tokens of the 'yes' type from a CPMM, the liquidity pool will have less 'yes' tokens and more 'no' tokens. This can be a risk for liquidity providers: If market participants continue buying 'yes' and 'yes' is the correct outcome, the liquidity pool ends up with worthless 'no' tokens. The phenomena of the liquidity pool loosing value as the price changes is also referred to as "impermanent loss" in the literature [10].

A variant of liquidity provision implemented by some AMMs is *concentrated liquidity* [11, 12]. In this variant, a liquidity providers provides a liquidity position that is only active for a specific price interval, i.e. for a specific interval of proportions of outcome tokens in the liquidity pool.

2.2.3 Market Resolution

For a prediction market to be useful, it must be possible to resolve the market, i.e. determine which of the traded outcomes turned out to be correct. Market resolution is not a trivial task, since there often is some level of ambiguity in the question that the market is predicting on. Furthermore, the entity which is resolving the market must be trusted. In centralised prediction markets, the entity resolving the market is the market operator. For example, the Iowa Electronics Market is operated by the university of Iowa [4].

In smart-contract based prediction markets such as Polymarket, the smart-contract that is underlying a market is typically resolved by an oracle. For example, the UMA oracle is used by Polymarket to resolve markets. The UMA oracle is an optimistic oracle, meaning that the UMA's resolution process is only triggered if someone makes a dispute. For most markets on Polymarket, the question is resolved by the prices for one of the outcomes tending to 1 USD, e.g. the market reaching consensus to price 'yes' at 1 USD. Any market participant can however dispute the resolution of the market. If a dispute is made, the holders of UMA tokens vote on the resolution of the market. Therefore, with such a decentralised resolution mechanism the trust in the market operator is replaced by the trust in the UMA token holders. As UMA tokens can be bought and sold anonymously, this exposes prediction markets to a risk of manipulation by UMA token holders.

2.3 Game Theoretic Properties

2.3.1 Insider Trading

One of the concerns that is sometimes raised when discussing prediction markets is insider trading. Insider trading is the trading of assets based on inside information that's not publicly available. For example, if a market participant knows from a non-public source that an asset's price is going to increase, or has the power to influence a real-world event that is important for the pricing of the asset, they could attempt to profit from this knowledge by trading with the asset. This problem also arises in markets other than prediction markets, for example in stock markets where insider trading is regulated by law. In prediction markets, insider trading is of particular concern because they are less regulated and often more anonymous than other markets.

What effects may insider trading have on prediction markets? One effect that insider trading could have on prediction markets is that it may affect their accuracy. There is however no consensus on whether insider trading might increase or decrease the accuracy of prediction markets. In the literature of the effects of insider trading some scholars argue that market participants with superior non-public information can increase the accuracy of a market by including this information in the pricing [13, 14]. On the other hand profits from insider trading incentivise to keep information secret, which could reduce the accuracy of a market [15]. Finally, insider trading could also detract market participants without superior non-public information from participating in the market, which could reduce the overall trading volume and thereby also the accuracy of prediction markets. In conclusion there is no consensus on whether insider trading might increase or decrease the accuracy of prediction markets.

2.3.2 Real-Money vs Play-Money

While prediction markets often trade with real money, there are also prediction markets that trade with play-money. For example, prediction markets on Polymarket are considered real-money markets as market participants trade with USDC (a "stable coin" pegged to the USD). On other platforms such as Manifold Markets[2] users bet with play-money that can not be converted to "real" money, i.e. it can not be converted to a currency that can be used to buy goods or services. There exists some debate on how the accuracy of real-money and play-money prediction markets compare, with some scholars having found real-money prediction markets to be more accurate [16] and others having found no significant difference in accuracy [17].

[2] https://manifold.markets/

2.3.3 Interpreting Prices as Predicted Probabilities

There exists some debate on whether prediction market prices should be interpreted as probabilities. As an initial example let's consider a market with two tokens, a 'yes' token and a 'no' token. The outcome which is resolved as correct will receive 1 USD. Let's further assume that the price is currently 0.60 USD for the 'yes' token and 0.40 USD for the 'no' token. Should we interpret this as the market participants predicting that the outcome is 60% likely to be 'yes' and 40% likely to be 'no'? If the true probability of the outcome being 'yes' is not 60%, then (assuming zero transaction fees) there exists a financial opportunity: If the true probability is higher than 60%, one can buy the 'yes' token for 0.60 USD and make profit in expectation. If the true probability is lower than 60%, one can buy the 'no' token for 0.40 USD and make profit in expectation. Therefore, an unbiased and expectation maximizing market participant would execute trades that push the price of an outcome token towards the true probability.

Wolfers et al. [18] make a more refined analysis of such a binary outcome market by examining a model in which traders have their wealth and their beliefs about the true probabilities drawn independently from a distribution. They assume that market participants maximize the expected value of a logarithmic utility function. Note that this implies a type of risk-aversion which will differentiate their model from other authors mentioned later. Finally, they assume that there are no trading fees (although transaction fees / costs are not explicitly discussed by the authors). With this model Wolfers et al. show that a market participant's demand x^* for an outcome token is dependent on their belief q about the true probability of the outcome being correct, and deduce

$$x^* = y \frac{q - \pi}{\pi(1 - \pi)}$$

where y is the wealth of the market participant and π is the price of the outcome token. By asserting that supply equals demand, Wolfers et al. show that the price of the outcome token is given by $\pi = \bar{q}$, i.e. the equilibrium price is the mean belief of the market participants.

This result breaks down when some of the assumptions are violated. While Wolfers et al. [18] were working on the model above in the years 2004-2006, there was a mutual influence with Manski [19] as well as with Gjerstad et al. [20]. Their analyses disagreed on the assumption of risk-neutrality vs. risk-aversion in particular. Manski [19] showed that under the assumption of risk-neutrality the equilibrium price is not necessarily the mean belief of the market participants, although it does give a bound on the mean belief.

In conclusion, the interpretation of prediction market prices depends on assumptions on beliefs, budgets, and risk preferences [18, 19, 20]. Under certain reasonable assumption the price of an outcome equals the mean belief of the market participants on the true probability. Furthermore, if the price of an outcome token is not equal to the real probability, then there exists a financial opportunity for market participants to make profit in expectation. We will therefore interpret prices as predicted probabilities in sections 3.3 and 3.4 in order to assess biases and accuracy of the markets.

2.3.4 Scoring Rules

For evaluating the accuracy of a prediction, so-called *scoring rules* can be used. These scoring rules were not specifically developed for prediction markets, but are widely used in this context for example by Hanson et al. [7, 21] who also laid the foundations for AMMs (see section 2.2.2). Consider the problem of scoring a reported probability distribution $\mathbf{p} = \{p_i\}_i$ over a set of disjoint events i. Let $\mathbf{o} = \{o_i\}_i$ be the actual binary outcomes drawn from the probability distribution $\mathbf{r} = \{r_i\}_i$, were $o_i \in \{0, 1\}$ and $0 \leq r_i \leq 1$. A **scoring rule** $s(\mathbf{p})$ assigns a score to the reported probability distribution $\mathbf{p}$. One scoring rule that is commonly used is the *quadratic scoring rule* or also called the **Brier score** [22]:

$$B(\mathbf{p}) = \frac{1}{N} \sum_{i=1}^{N} (p_i - o_i)^2$$

This is essentially the mean squared error of the prediction. It is a scoring rule and attains values between 0 and 1, where 0 indicates a perfect prediction. One property that such a scoring rule should have is that it should be *proper*. A scoring rule is proper if the reported score which maximizes the expected score is the probability distribution from which outcomes are drawn, i.e.

$$\mathbf{r} = \arg\max_{\mathbf{p}} \mathbb{E}[s(\mathbf{p})]$$

We can check that the Brier score is proper by computing

$$\begin{aligned}
\arg\max_{\mathbf{p}} \mathbb{E}[B(\mathbf{p})] &= \arg\max_{\mathbf{p}} \frac{1}{N} \sum_{i=1}^{N} (r_i(p_i - 1)^2 + (1 - r_i)(p_i)^2 \\
&= \arg\max_{\mathbf{p}} \frac{1}{N} \sum_{i=1}^{N} r_i + p_i^2 - 2r_i p_i \\
&= \mathbf{r},
\end{aligned}$$

where the last step obtained by calculating the derivative with respect to p_i and setting it to 0. This implies that the Brier score is proper and hence market participants who report their belief for the real probability distribution also maximize their expected Brier score.

Chapter 3
Data Analysis

3.1 Data Source

The data used in this chapter is scraped from prediction markets on the Polymarket platform. It combines several data sources:

1. The Polymarket subgraph. The code base for this subgraph is maintained by the Polymarket team on GitHub[1] and deployed on The Graph[2].
2. Polymarket's strapi API[3]. This offers data such as the market questions and market outcomes for resolved markets.
3. A polygon archive node via alchemy[4]. This complements some of the ambiguous and incomplete data retrieved from the subgraph. For example, the subgraph assigns transactions from the same polygon block the same timestamp, which omits the order in which the transactions were included in the block.

There are some limitations of this data: In particular, there is a small fraction of transactions which does not fall into the schema of the subgraph. For example, in rare cases a transaction trades on several markets, but only the part of the transaction which trades on one market is recorded in the subgraph.

3.2 General Statistics

The obtained dataset on Polymarket contains 2'054'717 Buy/Sell transactions across 9'026 markets. It further contains 174'934 liquidity pool additions and 157'137 liquidity pool removals. Users on Polymarket are also able to split USDC

[1] https://github.com/Polymarket/polymarket-subgraph
[2] https://thegraph.com/hosted-service/subgraph/polymarket/matic-markets-7
[3] https://strapi-matic.poly.market/
[4] https://www.alchemy.com/

into equal amounts of outcome tokens, which occurred 409'645 times, and merge outcome tokens into USDC, which occurred 182'713 times. Finally users redeemed USDC from resolved markets 423'855 times.

A histogram showing the transaction volume over time is displayed in figure 3.1. To understand the irregularity of the transaction volume, it should be considered that the trading volume can spike if one or a few markets attract particularly high activity: For example, several markets during and in the aftermath of the 2020 USD election attracted high trading volume.

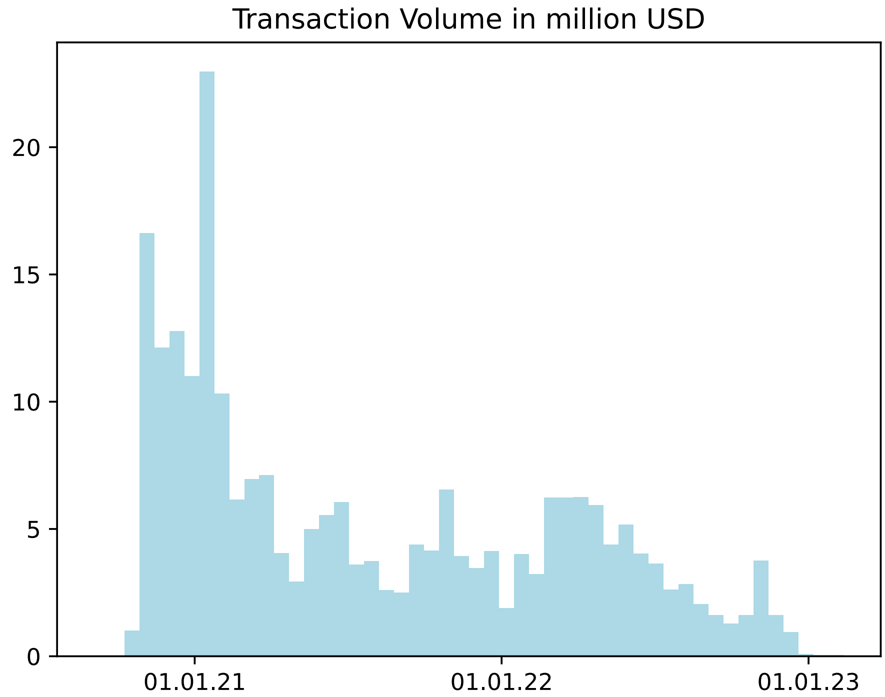

Figure 3.1: *Transaction volume over time.*

The distribution of the 9'026 market creations is presented in figure 3.2. To understand the irregularity of the number of newly created markets, it's important to recall that market creation is the only part of Polymarket which is not decentralized, but rather controlled by the Polymarket team. In the beginning of 2020 and into 2021, we saw a slow ramp-up of the number of markets created. Afterwards a period of high activity at the end of 2021 and beginning of 2022 followed, as the Polymarket team created many markets using repeating templates, for example about prices of crypto currencies, sport events and COVID-19. In the middle of 2022 Polymarket changed their market creation strategy and started to create markets with more unique questions, resulting in fewer new markets being created.

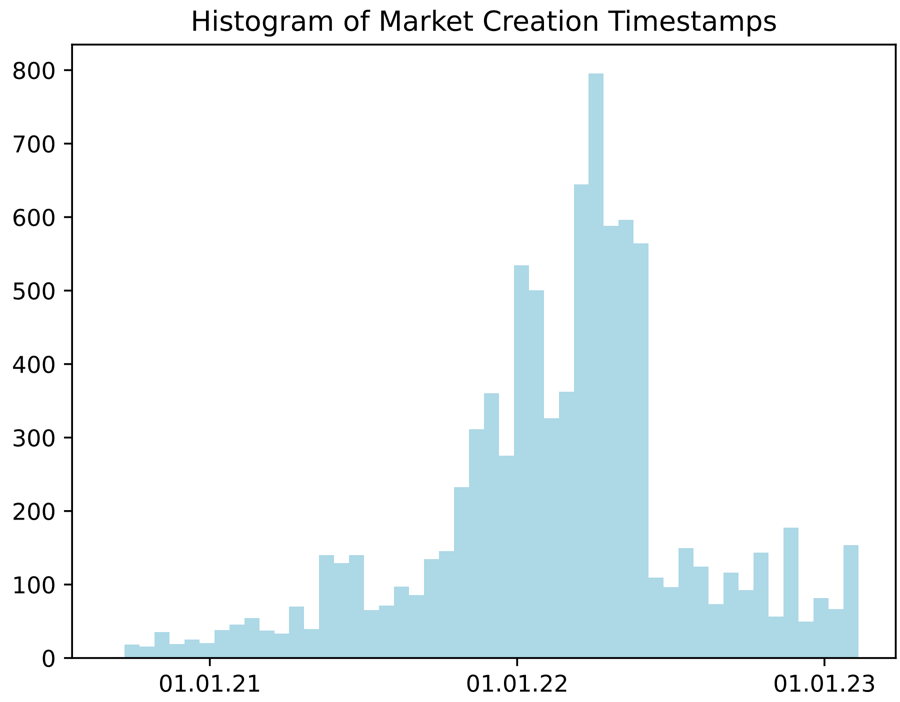

Figure 3.2: *Histogram of market creation timestamps.*

The Buy/Sell trades on Polymarket were executed by 34'691 users. Due to the anonymity of wallets it can not be determined how many of those users are unique personas, and how many are the same person utilizing multiple wallets. The 30 most active users by transaction volume are shown in figure 3.3. Overall, the 30 most active users account for 28.0% of all transaction volume, and the 100 most active users account for 50.5% of all transaction volume.

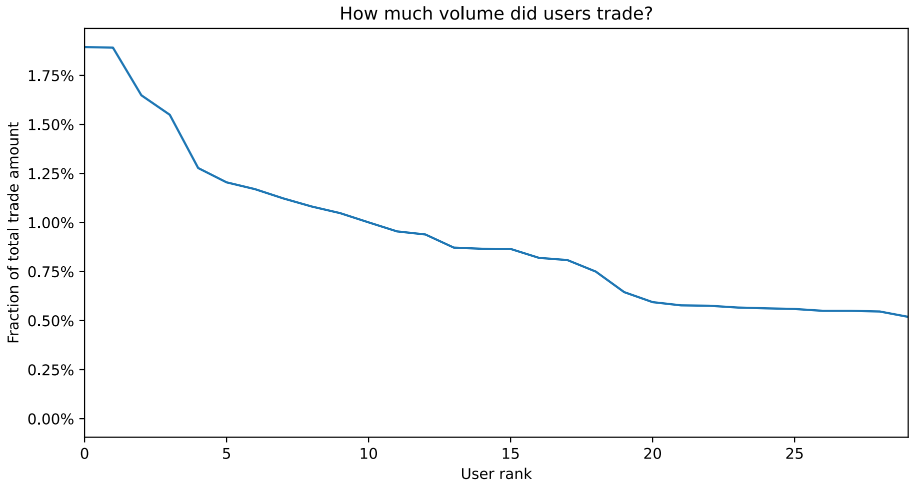

Figure 3.3: *The percentage of volume traded by the 30 most active users.*

3.3 Biases of Predictions

In order to assess biases of forecasting, some platform use a binary calibration plot [23]. Such a binary calibration graph plots a predicted probability versus the actual outcome. We adapt this approach and plot the prices of 'yes'/'no' tokens versus the fraction of these tokens for which the respective market was resolved 'yes'/'now'. If the prices would be unbiased, the graph would be a straight line with slope 1. We've plotted such a binary calibration graph for 'no' tokens in figure 3.4 and for 'yes' tokens in figure 3.5.

In figure 3.4 we see that for prices in the range of 0.2-0.8 that the markets resolve to 'no' more often than the price suggest. Put differently, in the price range 0.2-0.8 there is a bias of the 'no' token prices to be too low. In figure 3.5 we see the opposite bias for 'yes' tokens, i.e. the prices are too high for the 'yes' tokens in the range 0.2-0.8., albeit to a lesser extent than for 'no' tokens. Overall, this implies that on average 'yes' tokens have been overpriced, and 'no' tokens have been underpriced on Polymarket in the price range 0.2-0.8. While we can only speculate about the reasons for this bias, we note that the psychology literature has shown an "acquiescence bias" in some scenarios, which is the tendency to agree with agree-disagree questions [24].

When the price for 'yes' or 'no' tokens is close to 1, we see that there is a bias of prices being to low: Both in figure 3.4 and 3.5 the bin with the 0.80-0.90 price on the x axis corresponds to a fraction of accordingly resolved markets (resolved 'yes' for 'yes' tokens, 'no' for 'no' tokens) of over 90%. This implies that prices in the range 0.8-0.9 have on average been too low. Similarly, in the price range 0.10-.20, the fraction markets resolved 'yes'/'no' is below 10%, implying that prices

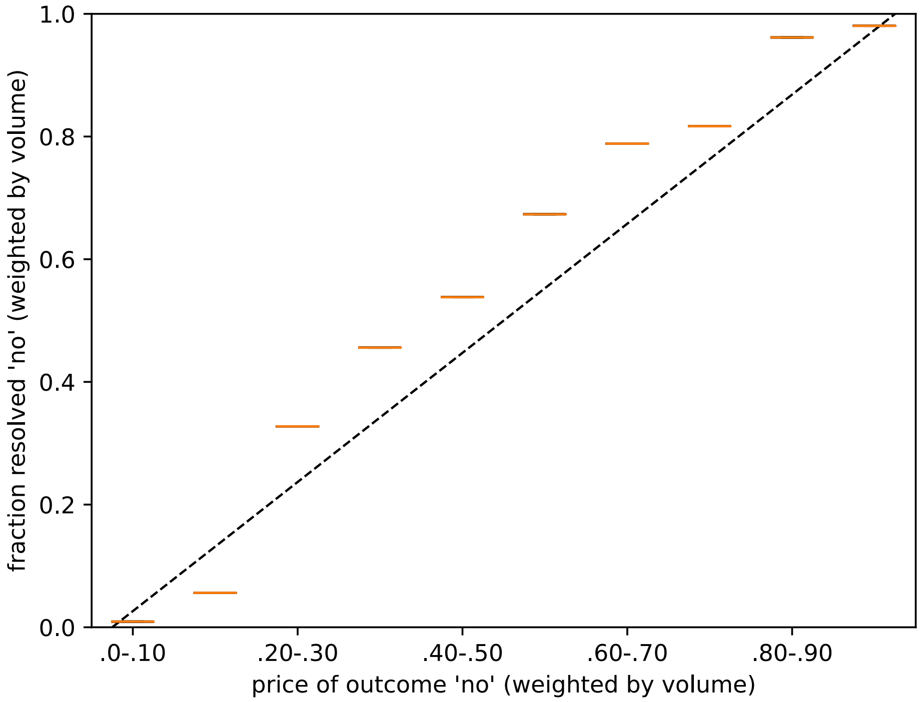

Figure 3.4: *A binary calibration graph for 'no' tokens: Every buy/sell transaction of a token falls into one of the buckets on the x axis depending on the transaction price. Per such bucket the fraction of tokens for which the corresponding market resolved to 'no' is plotted on the y axis.*

in this range have been too high. This could be due to various reasons: First, as the price of an outcome converges, the liquidity providers are incentivised to remove liquidity (as discussed in section 2.2.2) which makes it less favorable to execute transactions due to increasing price slippage. Secondly, if the price is close to 1, buying new tokens that agree with the market trend can only provide a small percentage of profit, which might be too small to justify the transaction costs and the opportunity costs of holding the tokens.

In conclusion, we have found that the prices of 'yes' and 'no' tokens on Polymarket have been biased as follows: When prices are in the range 0.2-0.8, the prices of 'no' tokens have on average been too low, and the prices of 'yes' tokens have on average been too high. When prices are below 0.2 or above 0.8, on average, there is a bias that the prices do not converge fast enough to 0 or 1, i.e. prices being to high for prices below 0.2 and too low for prices above 0.8.

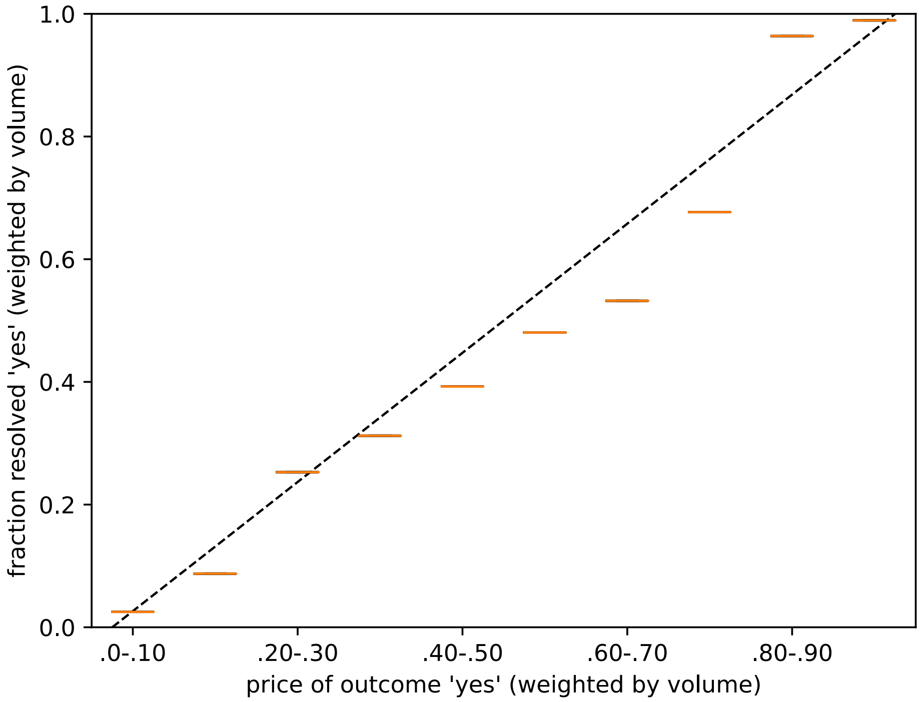

Figure 3.5: *A binary calibration graph for 'yes' tokens: Every buy/sell transaction of a token falls into one of the buckets on the x axis depending on the transaction price. Per such bucket the fraction of tokens for which the corresponding market resolved to 'yes' is plotted on the y axis.*

3.4 Accuracy of Predictions

As discussed in section 2.3.4 about scoring rules, one metric to assess the accuracy of predictions is the "Brier score". To recap, the Brier score assumes values between 0 and 1, where 0 indicates perfect accuracy. In this chapter we assess the accuracy of predictions on Polymarket in terms of Brier score relative to other dimensions such as time and transaction volume.

As a first perspective, we display the Brier score versus the market volume in figure 3.6. We see that the average Brier score of predictions on Polymarket is around 0.15, with no clear visible trend with respect to market volume. With high market volume the variance in the displayed figure increases. This can be explained by there being fewer markets for a bin with high market volume, and hence a small number of high-volume markets with especially low or high Brier score can impact one of the bins with high market volume.

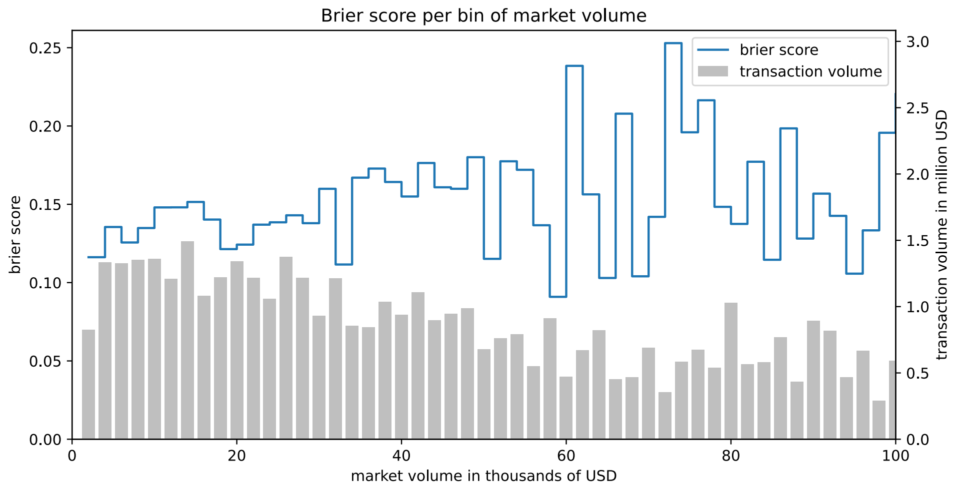

Figure 3.6: *The accuracy of predictions on Polymarket in terms of Brier score versus the market volume. Every buy/sell transaction of a token is interpreted as a prediction with the price indicating the predicted probability.*

Another interesting dimension to look at is the accuracy of predictions versus the date of the prediction, which is displayed in figure 3.7. The variance in this graph is very high, because at a given date one or a few markets might have dominated the trading volume. If such markets were particularly accurate or inaccurate, this is reflected in the graph. For example, around the end of 2020 there were markets about the 2020 US election on questions such as "Will Trump win [..]"[5] or questions on which party would win in the State of Georgia[6] and Pennsylvania[7]. Such markets saw a high trading value while the price of the correct outcome was already 0.8-0.9 USD, and hence the overall Brier score for predictions on Polymarket around this time is relatively low.

As trades occur over the lifetime of a prediction market, we would expect them to, on average, become more precise as time passes. For some markets the convergence to the correct market outcome happens quickly, as a response to some real-world event. In other cases it happens gradually, due to a series of new information or due to the absence of an event as a resolution date gets closer. By the time that information is available to resolve the market, we expect the price of the correct outcome to converge to 1 and hence the Brier score to fall. To test this empirically, we display the Brier score relative to time in figure 3.8, where the x-axis with the time represents the fraction of time between market creation and closure that has already passed. As expected, we can observe the Brier score

[5] https://polymarket.com/event/will-trump-win...
[6] https://polymarket.com/event/which-party-will-win-georgia...
[7] https://polymarket.com/event/which-party-will-win-pennsylvania...

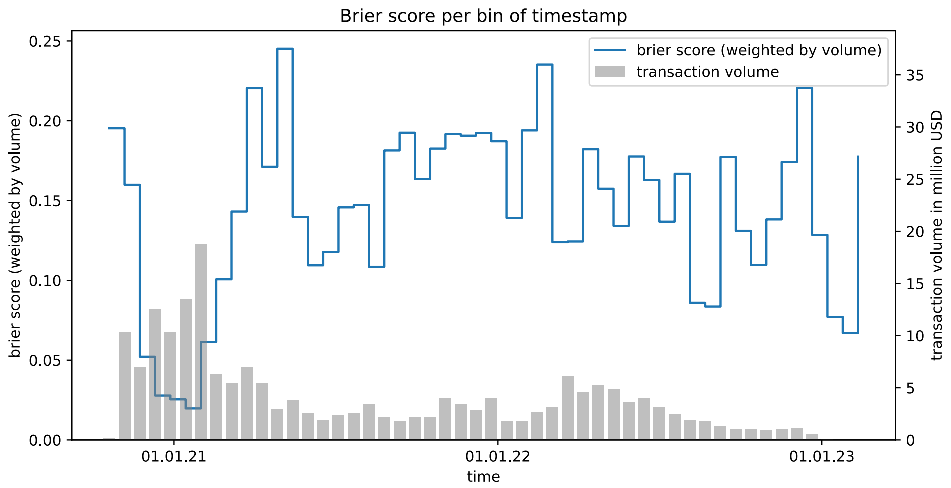

Figure 3.7: *The accuracy of predictions on Polymarket in terms of Brier score versus the date of the corresponding transaction. Every buy/sell transaction of a token is interpreted as a prediction with the price indicating the predicted probability.*

to drop as the closure time of the market comes close.

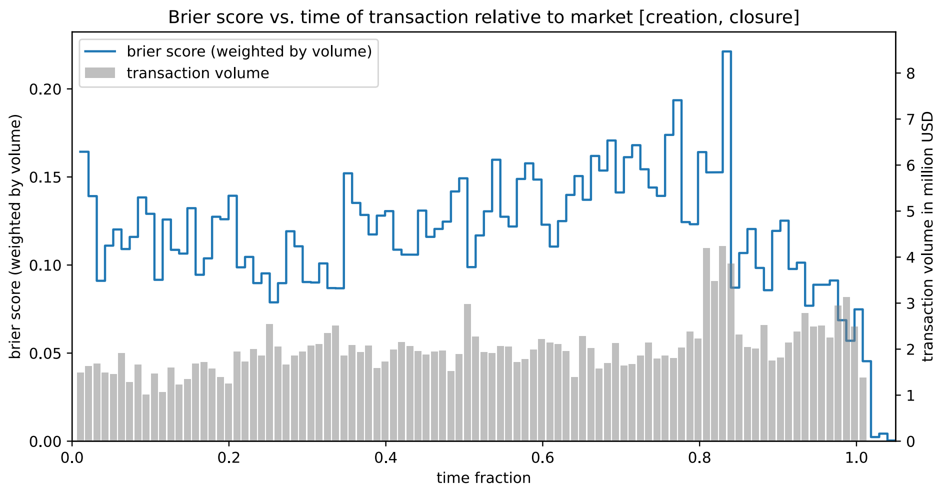

Figure 3.8: *The accuracy of predictions on Polymarket in terms of Brier score versus the fraction of time between the market creation and market closure that has already been passed when the corresponding transaction was executed. Every buy/sell transaction of a token is interpreted as a prediction with the price indicating the predicted probability.*

We investigate the relationship between the accuracy of predictions and the volume of the corresponding transaction in figure 3.9. As can be observed in the figure, there is a negative correlation between the Brier score and the volume of transaction. In other words: Historically, transaction on Polymarket with a volume of several thousand USD have, on average, been more accurate than transactions with smaller volume

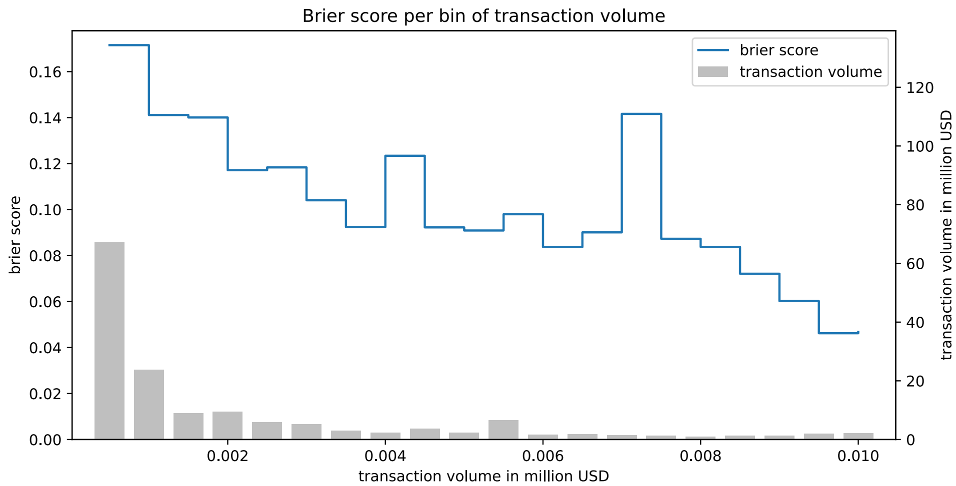

Figure 3.9: *The accuracy of predictions on Polymarket in terms of Brier score versus the volume of the corresponding transaction. Every buy/sell transaction of a token is interpreted as a prediction with the price indicating the predicted probability.*

3.5 Liquidity Provisioning

As described in section 2.2.2 about automated market makers (AMMs), many prediction markets are facilitated by AMMs. These rely on liquidity providers to fund the liquidity pool of tokens held by the AMM. In this section we will analyse some characteristics of the liquidity provisioning on Polymarket.

We begin by looking at the segmentation of liquidity providers: Are there a few dominant liquidity providers or are there many small liquidity providers? In figure 3.10 we can see the fraction of liquidity shares minted per user by the 30 largest liquidity providers. Recall that liquidity shares get minted as liquidity providers add funding to a liquidity pool, in order to track which user holds what fraction of the liquidity pool. In the figure we observe that there is one user who minted an extraordinarily large fraction of liquidity shares with circa 15%. Overall, the 10 largest liquidity providers minted 51% of all liquidity shares, the 30 largest liquidity providers minted 74% of all liquidity shares and the 100 largest liquidity providers minted 94%. A total of 3885 unique user IDs has added liquidity at least once. Similarly to the transactions in section 3.2, also with liquidity provisioning we can not deduce whether these user accounts belong to separate entities, or whether some entities anonymously controls multiple of the corresponding wallets.

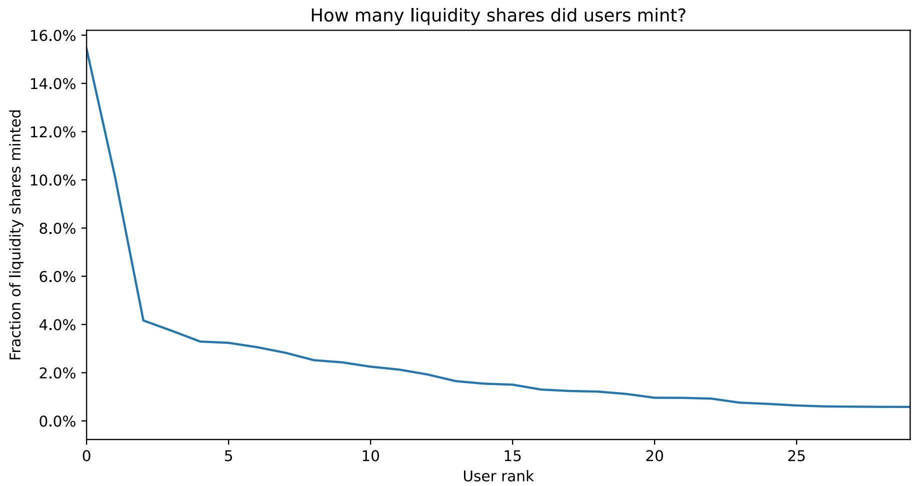

Figure 3.10: *The fraction of liquidity shares minted per user by the 30 largest liquidity providers.*

Another perspective from which to analyze the dominance of liquidity providers is this: First, we take the 10 largest liquidity providers across all markets, which corresponds to the ten largest liquidity providers in figure 3.10. For every market, we then calculate per market the fraction of liquidity shares held by these users over time. Finally, we average this across all markets (with all markets being weighted equally) and display the result in figure 3.11. As we can see the fixed set of 10 liquidity providers hold averaged across all markets about 25% of liquidity shares during the first half of a markets lifetime. After this the fraction drops to an average of circa 15%. This implies that the largest liquidity providers remove funds earlier, on average. An explanation for this could be that the largest liquidity providers might be more sophisticated and therefore remove their liquidity earlier as the risk of liquidity provisioning increases when the market converges (as explained in the "Liquidity Provision" part of section 2.2.2).

We've analysed the dominance of liquidity providers in terms of the fraction of liquidity shares they hold. But how many liquidity shares are there? This is displayed in figure 3.12. As in the previous figures, we calculated the number of liquidity shares per market over time, and average across the markets. We observe that on average, the number of liquidity shares increases over the first third of

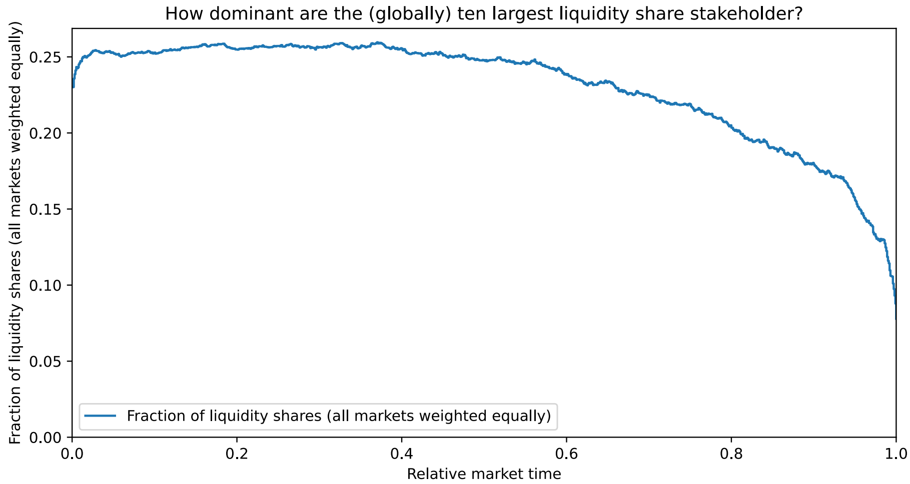

Figure 3.11: *The fraction of liquidity shares held by the 10 largest liquidity providers over time, averaged across all markets. The "largest" liquidity providers are determined globally, i.e. in terms of the liquidity shares they have minted across all markets.*

a market's lifespan, briefly plateaus, and then gradually declines. Shortly before the closure of a market the liquidity providers remove their shares. This can be explained by the fact that converging prices pose a risk to liquidity providers. We will use the characteristics of this curve in the next step, where we compare prediction accuracy vs. market liquidity.

The amount of liquidity provided to markets is a frequent concern among market participants, because the liquidity pool is the intermediary market participants may trade with. Therefore, too little liquidity may discourage market participation. Is this reflected in the prediction accuracy of markets? We investigate this in figure 3.13.

In this graphic, every buy/sell transaction of a token is interpreted as a prediction with the price indicating the predicted probability. These predictions are binned by the number of liquidity shares at the time of the transaction. Finally, based on our finding in figure 3.12, we only consider transactions that occurred at a time when 10%-80% of the market's lifetime has passed. This is to reduce the effect that predictions very early in a market or close to market closure land in different liquidity bins than earlier predictions, due to the average liquidity being lower at these times. In the figure we observe that empirically predictions in markets with liquidity roughly equivalent to 50'000 USD[8] have been substan-

[8]Adding 50'000 USD of liquidity to an empty liquidity pool gives this number of liquidity shares. For a non-empty pool the calculation is more complicated. However, given a certain state of the pool the amount of added funding and the number of received liquidity shares

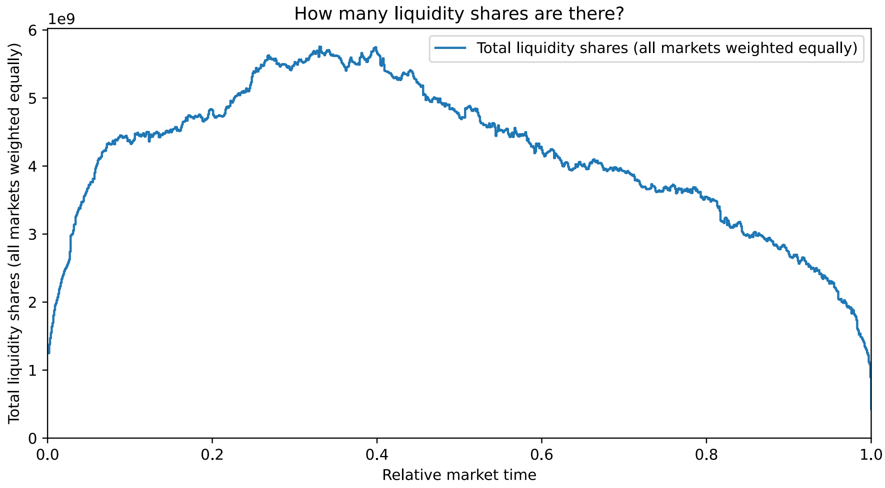

Figure 3.12: *The number of liquidity shares over time, averaged across all markets.*

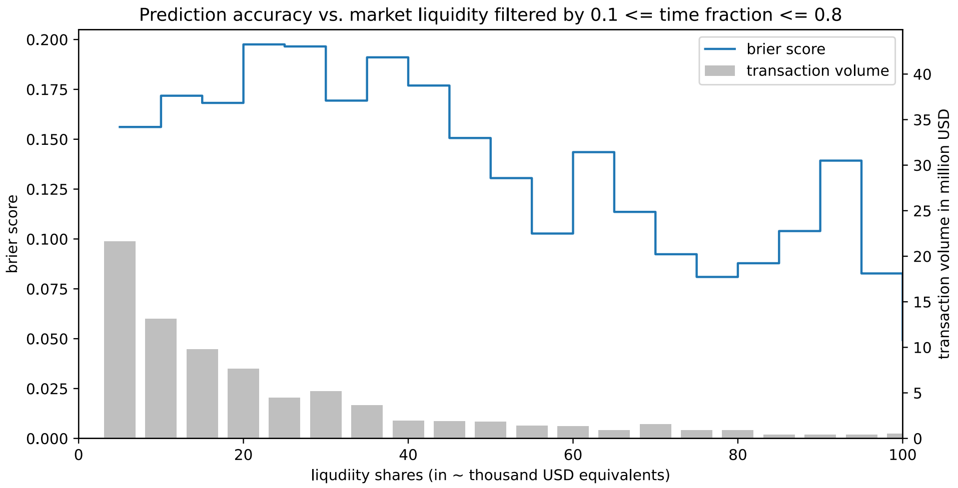

Figure 3.13: *The accuracy of predictions vs. the number of liquidity shares at the time of the transaction. Every buy/sell transaction of a token is interpreted as a prediction with the price indicating the predicted probability. These predictions are binned by the number of liquidity shares at the time of the transaction. We only consider transactions that occurred at a time when 10%-80% of the market's lifetime has passed in order to avoid some of the correlation between the amount of liquidity and the transaction time, which is visible in figure 3.12.*

remains proportional.

tially more accurate. Hence we do indeed observe in the data that markets with large liquidity have on average been more accurate.

Chapter 4

Discussion

4.1 Smooth Liquidity Provision

As we have seen in section 3.5, the average liquidity in markets drops as markets approach resolution. This can be explained by liquidity providers reducing their risk exposure as the prices converges. The problem with this outflow of liquidity is that markets with less liquidity are less accurate (as we saw in 3.4). In particular markets with converging prices are biased to not converge fast enough (as we saw in 3.3). In this section we propose an alternative automated market maker design that aims to mitigate the problem of reduced liquidity in converging prediction markets..

4.1.1 Idea

With a constant-product market maker (CPMM) it is rational for liquidity providers to reduce liquidity as the prices converge. If they do not do so and prices converge, they end up with a large number of tokens from the outcome(s) for which the price converges to 0 (see 2.2.2). In order to improve this situation, we would need to reduce the risk exposure of liquidity providers, so that they are not incentivised to withdraw liquidity as prices converge. Could it be effective to aim for zero risk exposure? Recall that prediction markets are a zero-sum game between buyers and sellers of outcome tokens and liquidity providers. Consider the example of a market where the price of an outcome token is known to be underpriced and therefore there are only traders who want to buy the token / sell other tokens for disjoint outcomes. If the traders are to make a profit in this scenario, some other market participant must lose money. In this case, it is the liquidity providers. Providing traders with an opportunity to make a buy transaction which is correcting an underpriced outcome token is necessary from a market design perspective. Therefore, we need to find a way to reduce the risk exposure of liquidity providers in order for liquidity providers to not be incentivised to withdraw liquidity as prices converge, while at the same time keeping liquidity providers exposed to a non-zero risk in order for traders to be able to make profitable trades.

In order to better understand the risk exposure of liquidity providers, we plot the number of tokens in the liquidity pool of a CPMM in figure 4.1. We assume that the market starts with a price of 0.50/0.50 for 'yes'/'no' and 100/100 tokens in the liquidity pool. We show how the number of tokens held by liquidity providers changes as the price converges to 1.00 for 'yes' while assuming that there are no liquidity additions or removals.

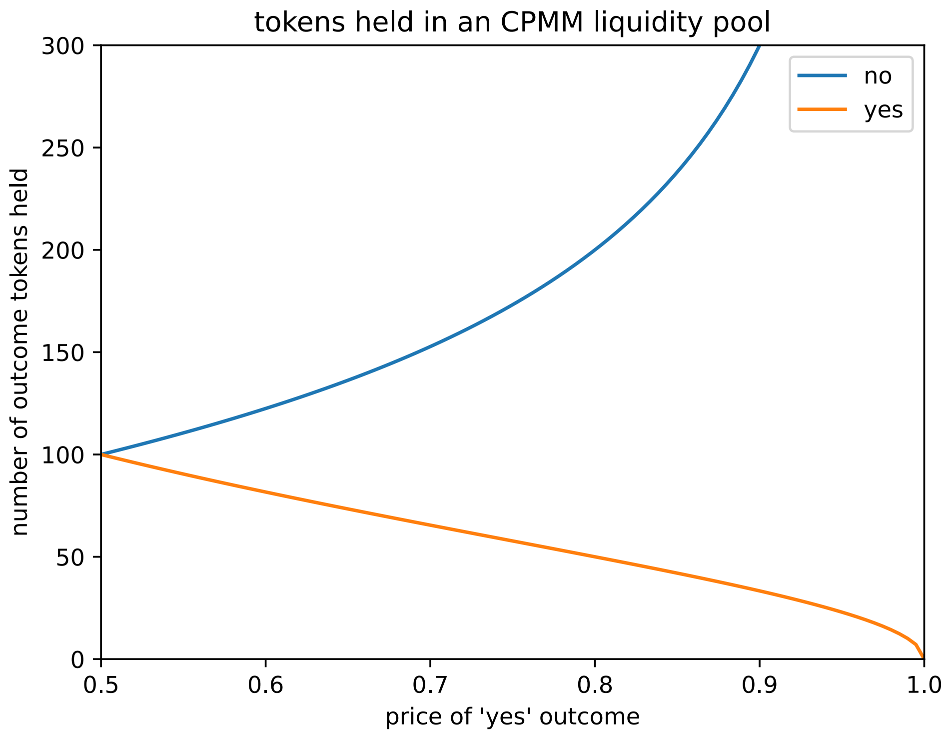

Figure 4.1: *A simulation of the number of tokens held by liquidity providers in a market with a constant-product market maker. The market starts with a price of 0.50/0.50 for 'yes'/'/no' and 100/100 tokens in the liquidity pool. We show how the number of tokens held by liquidity providers changes as the price converges to 1.00 for 'yes' while assuming that there are no liquidity additions or removals.*

We observe that as the price of the 'yes' tokens increases, the number of 'yes' tokens held by liquidity providers decreases and the number of 'no' tokens increases. This is the result of simultaneously fulfilling the formula for the CPMM

$$a \cdot b = k := a_0 \cdot b_0$$

,

where a, b, are the current number of tokens(e.g. a for 'yes' and b for 'no'), a_0, b_0 are the initial number of tokens, and k is the constant product, as well as fulfilling

$$p := \frac{a}{a+b},$$

where p is the current price. As we can see in the figure 4.1, as the price converges to 1.00, the liquidity providers end up with presumably worthless 'no' tokens, and hence are incentivised to remove their liquidity before the price converges. In order to solve this problem, we propose a new market maker design for which the number of valuable tokens held by liquidity providers don't converge to 0 as the price converges. We call this market maker design the smooth liquidity market maker.

4.1.2 Smooth Liquidity Market Maker

Consider the original CPMM formula

$$a \cdot b = k := a_0 \cdot b_0$$

where a, b, are the current number of tokens(e.g. a for 'yes' and b for 'no'), a_0, b_0 are the initial number of tokens, and k is the constant product. We modify this formula by 1) adding replacing k with a term that depends on $c(p)$ which we will refer to as the fraction of liquidity concentration, where p is the current price, and 2) on the left hand side subtract $c(p)a_0$ and $c(p)b_0$ from a and b respectively. This gives us the following formula:

$$(a - c(p)a_0) \cdot (b - c(p)b_0) = c(p)^2 \cdot a_0 \cdot b_0$$

The function used for $c(p)$ is adjustable and gives rise to a tradeoff between risk exposure of liquidity providers and incentives for traders (more on this in section 4.1.3). As an example, we propose

$$c(p) := 1 - |p - 0.5|$$

which is displayed in figure 4.2a. Simulating this smooth liquidity market maker with an example market gives rise to the liquidity pool displayed in 4.2b. As we can see, the number of valuable tokens held by liquidity providers converges to a non-zero value as the price converges to 1.00. This reduces the risk exposure of liquidity providers and hence reduces the incentive for liquidity providers to remove their liquidity as prices converge.

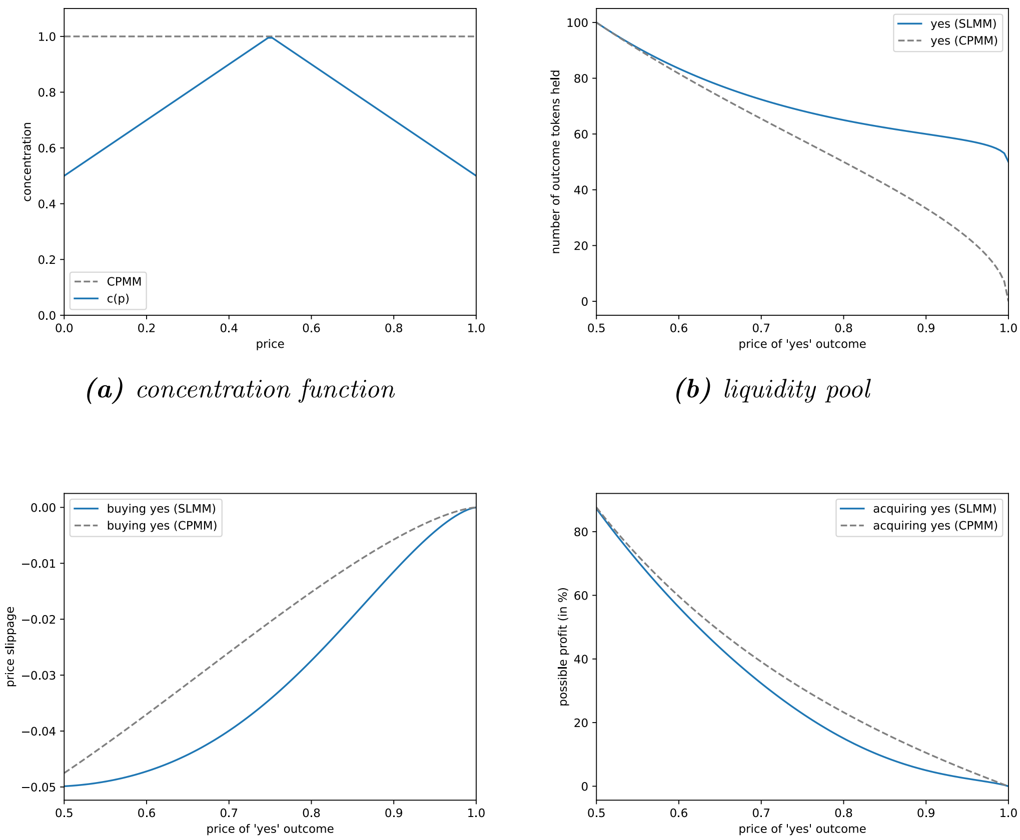

(a) *concentration function*

(b) *liquidity pool*

(c) *price slippage*

(d) *possible return on 'yes'*

Figure 4.2: *Example of a smooth liquidity market maker (SLMM) for a market with 'yes'/'no' outcomes. We assume an initial liquidity of 100/100 and no further liquidity additions or removals. The concentration function* $c(p)$ *in 4.2a for this example is* $c(p) = -2(p - 0.5)^2 + 1$*. For calculating the price slippage in 4.2c and the possible profit in 4.2d for investing in 'yes', we assume 2% liquidity fees and employ the tactic of splitting 10 USDC in 'yes'/'no' and selling 'no', which is more profitable than buying 'yes' directly.*

4.1.3 Trade-Offs

We have seen how the smooth liquidity market maker with the concentration function in figure 4.2a reduces the risk exposure of liquidity providers as visible in figure 4.2b. However, this reduced risk exposure comes at the cost of reduced incentives for traders. Concretely, the price slippage for a fixed liquidity pool and trade amount is increased, as is visible in figure 4.2c. As a consequence of this, the possible profit for traders on a fixed trade amount is reduced, as displayed in figure 4.2d.

In practical terms, if a trader observes e.g. price 0.70 USD and believes the asset should be priced at 0.75 USD, the absolute trade amount needed until the price reached 0.75 USD is reduced by the smooth liquidity market maker. This means that the trader can invest less capital to reach the valuation that the trader believes to be correct. In some instances this may lead a trader to not invest at all, as they might not wish to execute trades with such a small trade amount. In most cases this is a trade-off worth making, as it reduces the risk exposure of liquidity providers and hence reduces the incentive for liquidity providers to remove their liquidity as prices converge.

4.2 Response to "Five Open Questions about Prediction Markets" by J. Wolfers & E. Zitzewitz

In 2006 J. Wolfers and E. Zitzewitz published a paper titled "Five Open Questions about Prediction Markets" [25]. However, in 2006 this was only possible through the somewhat limited lense of 1) theory and 2) some niche markets such as university experiments, markets at companies, and sports betting. In this chapter we respond to three of their five questions on the basis of our data analysis of the diverse prediction markets on Polymarket.

Are markets well calibrated on small probabilities?

Concerned by evidence from the university of Iowa's election prediction markets, Wolfers and Zitzewitz [25] bring into question whether markets are well calibrated on events with small probabilities. They also bring forth potential reasons from psychology and risk-aversion for why this may not be the case.

Based on the empirical evidence form Polymarket, in particular section 3.3, we can answer this question: No, markets on Polymarket have not been well calibrated for small probabilities. In addition to the reasons from psychology and risk-aversion this may also be due to 1) A lack of liquidity for converging markets, as with Polymarket's AMMs liquidity providers operate at a loss as a market converges (see section 2.2.2), 2) transaction fees that outpace potential gains as the market converges, and 3) the opportunity cost of capital.

How to tradeoff interest and contractability?

The authors point out the difficulty of designing a market that is about an interesting topic yet has a sufficiently precise resolution criteria. Certainly, this has remained a challenge for prediction markets and hence great care is taken to precisely phrase the question and resolution sources. One tool that has been popularized since the work of Wolfers and Zitzewitz [25] to address this challenge are optimistic resolution oracles as described in section 2.2.3: This improves the tradeoff between interest and contractability by defaulting the resolution to the converged price, and only escalating to the oracle if the resolution is disputed.

Attracting uninformed traders

Wolfers and Zitzewitz [25] note that it could be challenging for prediction markets to attract uninformed traders, i.e. traders who operate at a loss so that informed traders may be incentivised by profits. They speculate on the reasons why uninformed traders may be motivated, such as hedging, overconfidence, and gambling. While we can not provide detailed insight into the motivation of uninformed trades on Polymarket, based on the significant trading volume on Polymarket on a wide range of topics we can answer Wolfers and Zitzewitz' question "Will prediction markets attract necessary uninformed trade" [25] positively.

Chapter 5

Conclusion

This practical project set out to understand the state of the art of prediction markets through theoretical and data analysis, as well as to improve the state of the art where possible. To this end we began by reviewing the theoretical background in chapter 2. In particular, we described the mechanisms of order books, automated market makers and market resolutions, as well as some game theoretic properties of prediction markets. From a theoretical perspective we already saw a potential problem with automated market makers for prediction markets in section 2.2.2: As a market converges, liquidity providers are incentivised to remove their liquidity, which may inhibit trading activity.

In a second step we scraped and analysed over 2 million transactions from the prediction markets on Polymarket. We found various biases in this historic data, and analysed the markets' accuracy against several other dimensions. We also performed a specialised analysis on the liquidity provisioning in these markets. In particular, we found that liquidity providers remove their liquidity as a market converges, and that markets become more biased and less accurate as they converge and as liquidity drops.

After observing the problem of liquidity provisioning for converging prediction markets in our scraped data as well as discussing the theoretical incentive structure that can lead to this problem, we proposed a new market making algorithm in section 4.1. This "Smooth Liquidity Market Maker" is designed to reduce the risk exposure of liquidity providers as a market converges. We showed that this reduces the incentive for liquidity providers to remove their liquidity, which promises to increase trading activity and improve the accuracy of prediction markets.

Bibliography

[1] J. Wolfers and E. Zitzewitz, "Prediction markets," *Journal of economic perspectives*, vol. 18, no. 2, pp. 107–126, 2004.

[2] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," *white paper*, vol. 3, no. 37, pp. 2–1, 2014.

[3] M. Schweitzer, "Mathematik und sportwetten-das haus gewinnt immer/eingereicht von michael schweitzer," Ph.D. dissertation, Universität Linz, 2020.

[4] P. Gomme, "The iowa electronic markets," *Economic Commentary*, no. 4/15/2003, 2003.

[5] B. Cowgill, J. Wolfers, and E. Zitzewitz, "Using prediction markets to track information flows: Evidence from google." *Amma*, vol. 3, 2009.

[6] B. Cowgill and E. Zitzewitz, "Corporate prediction markets: Evidence from google, ford, and firm x," *The Review of Economic Studies*, vol. 82, no. 4, pp. 1309–1341, 2015.

[7] R. Hanson, "Combinatorial information market design," *Information Systems Frontiers*, vol. 5, pp. 107–119, 2003.

[8] O. Guéant, "Optimal market making," *Applied Mathematical Finance*, vol. 24, no. 2, pp. 112–154, 2017.

[9] G. Angeris and T. Chitra, "Improved price oracles: Constant function market makers," in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 2020, pp. 80–91.

[10] L. Heimbach, E. Schertenleib, and R. Wattenhofer, "Risks and returns of uniswap v3 liquidity providers," *arXiv preprint arXiv:2205.08904*, 2022.

[11] H. Adams, N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson, "Uniswap v3 core," *Tech. rep., Uniswap, Tech. Rep.*, 2021.

[12] R. Fritsch, "Concentrated liquidity in automated market makers," in *Proceedings of the 2021 ACM CCS Workshop on Decentralized Finance and Security*, 2021, pp. 15–20.

[13] R. W. McGee, "Applying ethics to insider trading," *Journal of business ethics*, vol. 77, pp. 205–217, 2008.

[14] P.-J. Engelen and L. Van Liedekerke, "The ethics of insider trading revisited," *Journal of Business Ethics*, vol. 74, pp. 497–507, 2007.

[15] J. Surowiecki, "Capital gains," *The New Yorker*, 2005.

[16] S. Diemer, *Real-Money vs. Play-Money Forecasting Accuracy in Online Prediction Markets.* GRIN Verlag, 2013.

[17] E. Servan-Schreiber, J. Wolfers, D. M. Pennock, and B. Galebach, "Prediction markets: Does money matter?" *Electronic markets*, vol. 14, no. 3, pp. 243–251, 2004.

[18] J. Wolfers and E. Zitzewitz, "Interpreting prediction market prices as probabilities," 2006.

[19] C. F. Manski, "Interpreting the predictions of prediction markets," *economics letters*, vol. 91, no. 3, pp. 425–429, 2006.

[20] S. Gjerstad and M. Hall, "Risk aversion, beliefs, and prediction market equilibrium," *Economic Science Laboratory, University of Arizona*, 2005.

[21] R. Hanson, "Logarithmic markets coring rules for modular combinatorial information aggregation," *The Journal of Prediction Markets*, vol. 1, no. 1, pp. 3–15, 2007.

[22] G. W. Brier *et al.*, "Verification of forecasts expressed in terms of probability," *Monthly weather review*, vol. 78, no. 1, pp. 1–3, 1950.

[23] Metaculus. (2023) Metaculus: Track record. [Online]. Available: https://www.metaculus.com/questions/track-record/

[24] O. Kuru and J. Pasek, "Improving social media measurement in surveys: Avoiding acquiescence bias in facebook research," *Computers in Human Behavior*, vol. 57, pp. 82–92, 2016.

[25] J. Wolfers and E. Zitzewitz, "Five open questions about prediction markets," 2006.