Prediction Markets: Exploring Insights, Limitations, and Practical Applications

Introduction

Historically, prediction markets trace their roots back to early experiments in political and economic forecasting, such as the Iowa Electronic Markets in the late 1980's, allowing participants for the first time to trade contracts tied to the outcomes of specific events. Since then, these markets have evolved from experimental tools in academic settings to more sophisticated systems that drive decision-making in finance, politics, and beyond, offering forecasts in complex environments. Especially with the previous US elections, prediction markets have seen an increase in popularity, as their ability to aggregate diverse opinions has drawn significant attention from the media and the public.

At their core, prediction markets function by allowing participants to trade contracts based on the likelihood of certain events occurring. The prices are determined by the consensus of the traders, utilising the principle of crowd intelligence for predicting certain events as opposed to relying on just a single opinion. The prices of these contracts adjust in real-time, reflecting the collective beliefs of participants about the probability of an outcome. However, contrary to public markets, where bets are placed on the effect of an event on asset prices, prediction markets enable users to bet on the event itself.

A prediction market starts with a forecast, serving as the basis for constructing user scenarios. Within prediction markets, this is referred to as a condition. This condition can take the form of a question, statement, expected outcome, event, prediction, or any assertion about what might happen in the future. Every prediction has a set deadline to confirm whether it was accurate. On Polymarket for example, a designated oracle handles the task of deciding the outcome for each prediction.

This article will first outline the various types of prediction markets and touch upon their distinct concepts and market mechanisms. Furthermore, we will explore the platforms Kalshi and Polymarket, highlighting their differences, and later focus on the use cases of prediction markets. In the end, we will analyze accuracy compared to other forecasting methods.

Decentralized prediction markets

Most of the betting markets are decentralized. As we will explore further, there are different types of decentralized prediction markets. Decentralized prediction markets (DPMs) rely on two key components: smart contracts and oracles (Peterson et al., 2019). Smart contracts are self-executing programs where the terms of the agreement are embedded directly in the code, automatically carrying out the contract when specific conditions are met. Oracles act as intermediaries between the blockchain and the real world, supplying the external data needed for smart contracts to resolve outcomes.

In a DPM, users trade tokens representing potential outcomes of an event. For example, in predicting the winner of a sports match, participants might purchase tokens for "Team A wins" or "Team B wins." The token prices shift with market demand, reflecting the collective perception of each outcome's likelihood.

Once the event concludes, the oracle feeds the result to the smart contract, which distributes the pooled funds to those holding tokens for the correct outcome. This system ensures fairness and transparency, as the process is entirely governed by code, eliminating the need for a central authority.

Decentralized prediction markets (DPMs) offer several advantages over their traditional centralized counterparts:

- Transparency and Trust: Operating on public blockchains, DPMs provide a transparent and tamper-proof record of transactions and outcomes. This significantly reduces the risk of fraud or manipulation often seen in traditional betting platforms.

- Fair Odds: In DPMs, odds are determined entirely by market forces rather than being set by bookmakers, ensuring a fairer system that often results in better payouts for participants.

- Lower Costs: With no intermediaries involved, DPMs minimize fees and other costs associated with centralized betting or trading platforms. This makes participation more accessible and enhances the potential returns for users. If crypto is used for transactions (as we will see in the case of Polymarket), then gas fees apply.

Operational Models

As described above, prediction markets can be categorized into two distinct forms. However, each market is defined by the trading mechanisms, the way they facilitate forecasting, the level of decentralization and participant incentives. Below, the key concepts of prediction markets and their unique features are described.

1. Continuous Double Auction:

The continuous double action mechanism mimics traditional financial markets, where buyers and sellers are matched in real time. Participants place bids and offers for contracts tied to specific outcomes. The execution of the trade follows if a buyer's bid matches a seller's offer and the price of the transaction matches the consensus probability of the event. This continuous adjustment of prices reflects the collective beliefs of all market participants into a single, actionable probability. However, this mechanism can lead to challenges with liquidity for unpopular events, resulting in wide bid-ask spreads. Furthermore, lack of activity can also imply severely mispriced contracts. Still several prediction market platforms are based on CDA. One of them is the by the Commodity Futures Trading Commission regulated Kalshi, which will be touched upon later in this article.

2. Automated Market Makers (AMMs):

In contrast to CDA, Automated Market Makers (AMMs) are decentralized trading mechanisms. Meaning that assets can be exchanged without the use of traditional order books or direct counterparty matching. Instead, they rely on liquidity pools and mathematical algorithms, trying to facilitate trades by ensuring continuous market activity and liquidity. As described, AMMs don't match buyers and sellers directly; the process rather relies on a set of rules that are encoded into a smart contract, which again manages a reserve of tokens, the liquidity pool. These reserves consist of token pairs that, for example, are made up of a combination of ETH and USD, which both have a certain value. Now, whenever a trade is being executed, the ratio of tokens changes and therefore the price. In the long term, the prices adjust and reflect the market's view on a certain event. Additionally, overtime trading fees accumulate, which are later distributed amongst the LPs, delivering a reward for providing reserves for the liquidity pool. Even though that approach might theoretically imply continuous market liquidity, in reality, AMMs still face certain challenges. One of them is the slippage that occurs when a trader places an order in a rather small liquidity pool. This might push the price higher than expected, resulting in a less favorable trade. That phenomenon again discourages larger trades and therefore reduces market efficiency, especially in low liquidity pools. Other challenges, for example are high gas fees on the blockchain, impermanent loss, liquidity fragmentation by multiple liquidity pools, or dependance on initial liquidity.

3. Oracles:

Another key concept of prediction markets are oracles that are being utilized in decentralized markets. The main objective for oracles is to provide external data to resolve prediction market outcomes. By sourcing and verifying external information such as election results, sports outcomes, or financial metrics, oracles deliver the critical inputs for smart contracts to facilitate market settlement. The integrity and efficiency of a prediction market largely hinges on the quality of the data provided by oracles, as they inform the market's outcomes and ensure that participants can trust the system to be fair and transparent. Therefore, oracles are the backbone of prediction markets. Oracles can be generally divided into centralized and decentralized types. In comparison to decentralized oracles, centralized oracles do not aggregate data from multiple sources bur rather rely on a single trusted data source, making them faster, but more vulnerable to manipulation. Since real-world events can pose unique challenges, for example the slow counting of votes during the 2020 US presidential election, oracles have to navigate these complexities in order to provide accurate data to the markets. Their integration ensures that prediction markets operate transparently and efficiently.

Kalshi – the leading centralized prediction market

Compared to Polymarket Kalshi is one of the very few centralized prediction markets. Unlike many other firms, it obtained the licence of a fully regulated financial exchange, allowing it to operate under the oversight of the Commodity Futures Trading Commission (CFTC). Therefore, in 2020, Kalshi became the first federally regulated exchange dedicated exclusively to trading event contracts. Since then, the New York based startup, which is backed by Sequoia, Y Combinator, Charles Schwab, Henry Kravis, Neo, SV Angel, and others, has seen steady growth, having now just recently reported a trade- volume of over $1bn.

In contrast to Polymarket Kalshi is centralized and therefore stronger resembles traditional stock exchanges. It utilizes the Continuous Double Auction mechanism, meaning that buyers and sellers are matched in real time. Furthermore, Kalshi oversees all aspects of trading on its platform, including order matching via CDA, contract settlement and regulatory compliance like Know Your Customer and Anti Money Laundering procedures. Similar to other prediction markets contracts are priced from $0.01 to $0.99, representing the market's perceived probability (1% to 99%) of an event occurring and contract settles at either $1 or $0 dependent on the events outcome. Contracts are offered across severall categories including markets for Politics, Economics, Weather, Crypto or Finance. Similiar to traditional exchanges, Kalshi collects transaction fees based on trade volume. Compared to financial markets, the platform does not allow for leverage.

This regulatory framework provides users with legal protection, enhancing market integrity through the prevention of fraud and manipulation. Especially, when it comes to contract resolution Kalshi is designed to be transparent and reliable, having clear definitions in their event contracts and further providing a structured process in case of disputes. Additionally, Kalshi is user-friendlier than Polymarket allowing for more accessibility for inexperienced users, that might be unfamiliar with the technological intricacies of blockchain technology. However, Kalshi therefore lacks in areas that decentralized markets excel in. They don't provide the same anonymity as markets built on blockchain technology and are sometimes limited by jurisdiction, facing regulatory barriers. This leads to less variety of offered markets, possibly limiting flexibility and innovation for users. Furthermore, Kalshi holds complete control over the traders' funds, meaning users rely on the platform for fund management and withdrawals.

Polymarket – the leading decentralized betting market

Polymarket is one of the top decentralized prediction market platforms. Built on the Polygon network, a sidechain of Ethereum, it utilizes the Universal Market Access (UMA) oracle to function. Polymarket does not have U.S. regulatory approval, leaving users without formal legal protections. While its decentralized nature may attract those experienced with DeFi, its legal status is unclear, particularly in the United States. The Commodity Futures Trading Commission (CFTC) imposed a $1.4m fine on Polymarket and required it to close its markets and refund users, citing a failure to register with the regulator.

Polymarket lets users place bets on a wide range of topics, including political events like the U.S. presidential election, sports events such as the 2024 Olympics, and cryptocurrency price movements. Political betting dominates Polymarket, accounting for approximately 70% of its activity, with individual bets varying from $1,000 to $100,000. In June 2024, the platform saw 29,432 new users and 3,590 daily active traders adjusting their positions. Open interest exceeded $46m, while daily trading volume reached $4.6m.

Due to the decentralized nature, to trade on Polymarket, one needs to have a wallet (since all the transactions are crypto-based). The wallet is generated using a dedicated smart contract factory and operates as a 1/1 multisig, giving the user complete control over it.

As mentioned in the introduction, a prediction market begins with a forecast, which has a certain end date. A typical Polymarket trading market can be divided into 3 main parts:

1. Creation of a market, where a new prediction (or condition) is introduced with a defined timeframe.

2. Users can then select the outcome they think is most probable and buy/sell "yes" or "no" contracts in case the market states a condition where something will or will not happen. An example could be: "will Bitcoin reach 100k before 2025?" Another possibility is deciding 1 out of all the listed options in case the question has more than 2 outcomes. An example could be "Which will be the highest-grossing movie in 2024?"

3. Once the market closes, the oracle identifies the correct outcome, determining which users were right and which were not.

On Polymarket, all event outcomes are tokenized. This means whenever you trade on Polymarket, you are essentially buying a share token, which is fully collateralized with the asset used to buy it. As with Khalshi, the price of a token indicated the perceived probability of the event occurring.
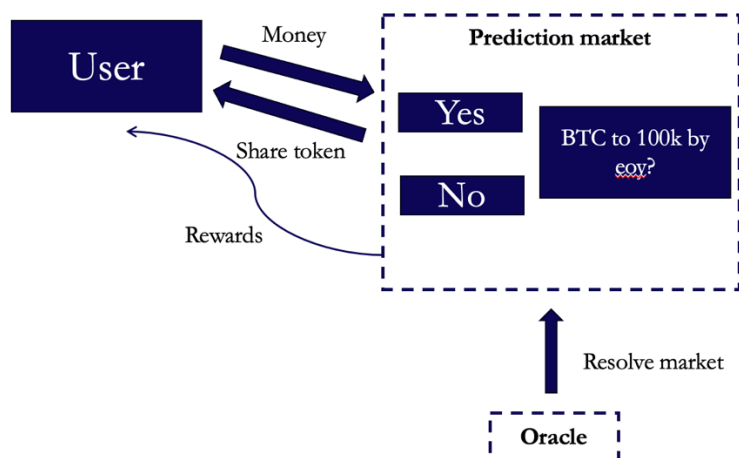
Share tokens are built on ERC-1155, leveraging the Gnosis Conditional Tokens Framework (CTF). CTF can support up to 256 outcomes per event. In this system, every prediction is uniquely identified by a conditionalId. This conditionalId is generated as a hash of three key components:

1. Oracle: The address of the oracle responsible for determining the event's outcome, ensuring that only the designated oracle can finalize the prediction.

2. QuestionId: A unique identifier for the prediction, defined by its creator.

3. OutcomeSlotCount: The total number of potential outcomes associated with the prediction.

This is how the mechanism works:

When the outcome of a market is determined, the market can be "resolved" or finalized permanently. Resolution follows the market's predefined rules, detailed in its order book. Upon resolution, holders of winning shares receive $1 per share, while losing shares become worthless, and further trading of shares is no longer allowed. To resolve a market, an outcome must first be "proposed," which requires posting a bond in USDC.e. This bond is forfeited if the proposal is deemed incorrect. If the proposed outcome is verified as correct, the proposer is rewarded for their contribution. If a proposal is deemed incorrect or disputed, the bond is forfeited, encouraging participants to act in good faith. This resolution mechanism ensures fairness, accuracy, and accountability while maintaining market integrity. At the resolution date, the correct outcome has a payoff of 1$. This means that the PnL for whatever buy price is: $PnL(\%) = \frac{Buy\ price - 1\$}{1\$}$.

Applications of Prediction Markets

Over the years prediction markets have evolved from academic experiments like the Iowa Electronic Markets to publicly accessible platforms like Polymarket and Kalshi, offering event-based trading to the public. This evolution has also increased the scope of their applications, which range from forecasting election outcomes, predicting economic indicators, estimating sales figures, anticipating technological advancements, to gauging public health developments during pandemics. Other use cases even include managing corporate project timelines, hedging event-specific risks, or evaluating policy impacts. Below, the theoretical applications will be highlighted, and their practical viability and limitations will be discussed.

Information Aggregation and Decision Support:

Prediction markets have developed into an interesting tool for information aggregation and decision support in many different areas. On Polymarket and Kalshi, political events have especially seen a lot of public interest. In that domain, markets predict the outcomes of elections, policy decisions, legislative developments and geopolitical events, therefore offering valuable insight into future political landscapes. Also, in the domain of financial markets and for predictions of economic indicators, prediction markets can be considered for making decisions. They provide insights on macroeconomic data releases such as unemployment rates, inflation, or GDP figures.

Furthermore, it is possible to deduct market sentiment by evaluating predictions for central bank actions. However, prediction markets have also been utilized for internal forecasting. For example, Google launched an internal prediction market on Google Cloud to leverage collective employee wisdom, combining that with machine learning models, to improve predictive analytics. Another area, in which prediction markets can provide insights, is public health and social trends. Markets, for example, forecast disease spread, societal shifts, or technological adoption. This can aid governments in policy decisions and provide guidance on emerging challenges.

Even though these predictions provide valuable insights by utilizing market efficiency and collective knowledge, they are far from accurate forecasts. Several factors contribute to their limitations. One of the main issues of prediction markets is their thin liquidity, making markets less efficient and susceptible to manipulation. Furthermore, they reflect several behavioural biases, such as overconfidence bias, anchoring effect, and herd behaviour. Additionally, the prediction accuracy is strongly impacted by the market's participants. Nevertheless, prediction markets serve as a viable tool to gain market insights and deliver forecasts for future events. To further determine the reliability of their predictions, we will later explore different research to evaluate the accuracy of prediction markets.

Trading and Risk Management:

In addition to providing forecasts, prediction markets open up new opportunities for investment. They provide possibilities to profit from correctly predicting event outcomes and therefore offer ways of investing very granularly that traditional financial markets would not allow for. That very nature of prediction markets theoretically also delivers a very direct approach to manage risks. Companies and investors could mitigate potential losses by leveraging prediction markets to hedges against adverse outcomes tied to specific events. This would allow them to offset risks associated with uncertainties such as regulatory shifts, macroeconomic changes, or geopolitical developments. For example, a company launching a new product could hedge against the risk of delayed approval by taking position in a suitable contract.

However, prediction markets are currently far from suitable for investment or hedging on a larger scale. The liquidity constraints make it very difficult to enter or exit positions at favourable prices. This reduces the effectiveness of hedging, especially for trades with higher volume. Furthermore, traditional financial markets offer alternative instruments that serve similar purposes. Even though their scope might not be as granular, they are currently much more accessible and efficient in hedging positions.

Accuracy of prediction markets

As mentioned above, prediction markets have emerged as a compelling tool for forecasting events. Their ability to aggregate diverse information across various domains, has been highly appreciated. Some even argue that, by leveraging the collective wisdom of participants, who have financial incentives to be accurate, prediction markets provide a better and more accurate estimation of certain events than other prediction methods. Even though their unique concept and their increase in popularity cannot be denied, their accuracy relative to polls, expert opinions, and statistical models, remains a subject of active debate. In this section, we explore evidence from various research studies to assess the accuracy of prediction markets and provide a balanced perspective on their utility.

To understand the research on the accuracy of prediction markets, the concepts of the Brier Score and the Brier Skill Score are critical:

Brier Score (BS): Measures the accuracy of probabilistic predictions by computing the average squared difference between predicted probabilities and actual outcomes. Lower scores indicate better predictions.

$$BS = \frac{1}{N}\sum_{i=1}^{N}(f_i - o_i)^2$$

Brier Skill Score (BSS): Compares the accuracy of a prediction model (BS $_{forecast}$) to a baseline model (BS $_{baseline}$). A higher BSS indicates better performance relative to the baseline.

$$BSS = 1 - \frac{BS_{forecast}}{BS_{baseline}}$$

When looking at the evidence that research provides, the results are mixed. For example, Jason Dana et al. (2023) conducted a study where they compared prediction markets with self-reported probabilities. Those probabilities were reported by over 500 participants (70% with postgraduate degrees), covering 113 geopolitical questions and later compared to the prices in an artificially created prediction market. The study concluded that aggregated self-reported beliefs were at least as accurate as prediction market prices, as measured by comparing the two Brier Scores. A study that shows slightly different results is prediction market accuracy in the long run, by Berg et al. (2008). In this case, the Iowa Electronic Markets demonstrated that prediction markets outperform election polls in the long run. Another study (Clay, 2022) that involves the CFTC regulated Kalshi prediction markets, evaluated datapoints of 8.746 finalized markets of the platform. The goal was to determine the predictive ability of Kalshi's markets compared to the base event rate. As expected, the markets showed a positive Brier Skill Score that increases with maturity to market closure. Even though that underlines the predictive ability, the study's results don't make prediction markets more accurate than polls, experts' opinions or statistical models.

In general, it can be concluded that prediction markets are a very interesting method to evaluate the likelihood of future events and to gauge market sentiment; however, it cannot be clearly stated that they are more accurate than other evaluation methods of future events.

Conclusion

We found that prediction markets offer a novel approach to forecasting by aggregating collective intelligence and financial incentives, yet they are not without their flaws. Low liquidity, susceptibility to manipulation, regulatory barriers, and participant biases can limit their effectiveness and accuracy. While they excel in providing insights into political, economic, and corporate outcomes, their performance compared to traditional methods like polls and models remains mixed. Despite these limitations, prediction markets demonstrate considerable potential for improving decision-making and understanding market sentiment, making them a valuable, albeit imperfect, tool for navigating uncertainty.

References

[1] Schrieber, Jared, "The application of prediction markets to business", 2004

[2] Dana, Jason, "Are markets more accurate than polls?", 2023

[3] Berg et. al, "Prediction Market Accuracy in the Long Run" 2008

[4] Clay, Russ, "Calibration and Skill of the Kalshi Prediction Markets", 2022

[5]   Jain, Anchit, "The Rise of Blockchain-Based Prediction Markets", 2024

[6] Chen, Yiling, et al. "Information markets vs. opinion pools: An empirical comparison", 2005.