Crypto Prediction Markets:

Collective Wisdom or a Hunting Ground for the Few?

About

Metanomia is a think tank focused on analyzing transformations in the cryptocurrency market and related technologies.

Researcher

Seyeon Jo

Corresponding Author

Taemin Oh

Contents

01	**Introduction**	3
02	**From Gambling to Collective Intelligence**	6
03	**Anatomy of Profit Strategies: Where Do Market Inefficiencies Come From?**	10
	The Technical Structure of Prediction Markets: Split, Merge, and Convert	
	Three Pathways of Logical Arbitrage	
	Who Took How Much?	
	Section Summary: Inefficiencies That Will Fade With Market Maturity, and the Questions That Remain	
04	**The Decay of Alpha and the Transformation of Collective Intelligence**	17
	The Decay of Alpha: Closing Profit Windows	
	What Remains: Three Types of Information Advantage	
	On-Chain Transparency: Insiders Are Surprisingly Visible	
	The Precedent of Stock Markets	
05	**Institutional Design to Safeguard Prediction Market Integrity**	22
06	**Conclusion: The Paradox of Collective Intelligence**	24
	References	26

01

Introduction

Prediction markets expanded dramatically following the 2024 U.S. presidential election. Polymarket recorded approximately $9 billion in cumulative trading volume in 2024,[1] and in 2025, the combined annual trading volume across major platforms rose to $44 billion.[2] This momentum continued into 2026: according to data from the Pew Research Center, the combined monthly trading volume of Kalshi and Polymarket reached approximately $24 billion in April alone.[3] In just over two years, prediction markets have expanded at a remarkable pace.

The decisive catalyst for this growth was the 2024 presidential election. While major public opinion polls predicted a closely contested race between Kamala Harris and Donald J. Trump, Polymarket prices consistently indicated a Trump lead in the period leading up to the election—and ultimately proved correct. As Polymarket's prediction aligned with the actual election outcome, the long-standing proposition that markets involving financial stakes can aggregate dispersed information into prices gained broad public credibility.

Of course, the growth of prediction markets during this period cannot be explained by accuracy alone. Changes in the U.S. regulatory environment, integration with major platforms, the expansion of sportsevent contracts, increased media exposure, and growing stablecoin liquidity all contributed. Nevertheless, the 2024 presidential election clearly served as the turning point that gave prediction markets a self-narrative as accurate information markets and brought them into the mainstream. The argument advanced earlier by economists, including Nobel laureate Kenneth Arrow, that prediction markets perform an important information-aggregation function and therefore warranted regulatory liberalization regained attention in light of the actual performance of these markets.[4]

1) Polymarket's cumulative trading volume in 2024 reached approximately $9 billion, with monthly trading volume reaching a then-record $2.63 billion in November. *The Block*, "Polymarket's Huge Year: $9 Billion in Volume and 314,000 Active Traders Redefine Prediction Markets," January 3, 2025, https://www.theblock.co/post/333050/polymarkets-huge-year-9-billion-in-volume-and-314000-active-traders-redefine-prediction-markets.

2) In 2025, the total annual trading volume of major prediction markets—including Polymarket ($21.5 billion) and Kalshi ($17.1 billion)—was estimated at approximately $44 billion. This represented roughly a fivefold increase from approximately $8.8 billion in 2024 across the same group of platforms. The $9 billion figure in footnote 1 refers to Polymarket alone; therefore, the two figures cover different sets of platforms. Boaz Sobrado, "How Prediction Markets Actually Grew in 2025," *Forbes*, December 16, 2025, https://www.forbes.com/sites/boazsobrado/2025/12/16/how-prediction-markets-actually-grew-in-2025/.

3) The combined monthly trading volume of Kalshi and Polymarket grew from less than $5 billion in September 2025 to approximately $24 billion in April 2026. Pew Research Center, "Trading Volume on Prediction Markets Has Soared in Recent Months," May 27, 2026, https://www.pewresearch.org/short-reads/2026/05/27/trading-volume-on-prediction-markets-has-soared-in-recent-months/.

4) Kenneth J. Arrow et al., "The Promise of Prediction Markets," *Science* 320, no. 5878 (2008): 877–878, https://doi.org/10.1126/science.1157679.

Following this turning point, prediction markets began to shed their stigma as little more than a sophisticated form of gambling. Their ability to aggregate dispersed information (information aggregation), display the probability of events through real-time prices (price discovery), and provide a means of offsetting risks associated with exposure to specific events (insurance effect) has strengthened the view that prediction markets should be regarded not merely as betting platforms, but as public-good infrastructure.

Moves by Big Tech have reinforced this shift. Following Robinhood's integration of prediction-market functionality into its platform and Coinbase's expansion of prediction-market features, Meta began developing Arena, its own prediction-market app.[5] Meta's entry is particularly symbolic. The company previously challenged the fiat monetary order through Libra, only to wind down the project amid congressional hearings and regulatory pressure. The fact that Meta has now entered a market that is still classified and prohibited as gambling in many countries, including South Korea and Japan, suggests that prediction markets have already become an increasingly difficult trend to ignore.

However, this report focuses not on the fanfare surrounding this growth, but on the paradox beneath it. The paradox is not a simple binary question of whether prediction markets constitute "wisdom of the crowd." Rather, prediction markets operate on two distinct levels. One is the information-aggregation layer, concerning how closely market prices approximate actual outcomes. The other is the distributive layer, concerning who has access to opportunities to participate in the market and who captures the resulting gains.

The reality revealed by on-chain data is clear. At the information-aggregation layer, prediction markets perform successfully: as more participants and algorithmic bots provide liquidity, market prices become more accurate than opinion polls. At the distributive layer, however, stark inequality persists. According to an on-chain analysis that aggregated the net profits and losses of approximately 1.73 million Polymarket addresses, more than 70% of total net profits were concentrated among fewer than 0.04% of addresses.[6]

The core thesis of this report, therefore, is not that prediction markets fail to embody the wisdom of the crowd. Rather, prediction markets can aggregate dispersed information into prices and produce signals that closely approximate actual outcomes. Yet the profits and sustained information advantages generated by this process may become concentrated among a small minority.

This raises two fundamental questions. Is the concentration of profits a necessary cost of producing accurate market prices, or does it indicate a market failure? And where should we draw the line between legitimate compensation for an information advantage and rents that should be institutionally constrained?

---

5) Robinhood launched its Prediction Markets Hub in March 2025, while Coinbase announced plans to expand its prediction market offerings in late 2025 and began offering a Kalshi-integrated service to U.S. users in early 2026. Meta was reported to be developing a prediction market app called *Arena* in June 2026. Robinhood, "Robinhood Launches Prediction Markets Hub," March 17, 2025, https://robinhood.com/us/en/newsroom/robinhood-prediction-markets-hub/ ; Coinbase, "System Update: The Future of Finance Is on Coinbase," December 17, 2025, https://www.coinbase.com/blog/system-update-the-future-of-finance-is-on-coinbase; "Coinbase (COIN) and Kalshi Team for Prediction Market for Crypto Exchange's U.S. Customers," *CoinDesk*, January 27, 2026, https://www.coindesk.com/markets/2026/01/27/coinbase-rolls-out-prediction-market-to-u-s-customers; NPR, "Meta Plans to Release AI-Powered Prediction Market App, Documents Show," June 24, 2026, https://www.npr.org/2026/06/24/nx-s1-5869486/meta-prediction-market-app-ai.

6) An on-chain analysis of approximately 1.73 million Polymarket addresses found that more than 70% of total net profits were concentrated among fewer than 0.04% of addresses. However, because the analysis is conducted at the address level, a single user may control multiple addresses, and its detailed methodology has not been academically validated. This report therefore uses the finding as a supplementary indicator of profit concentration. KuCoin, "Polymarket Profitability Crisis: 70% of Users Lose Money, Profits Concentrated among 0.04%," December 29, 2025, https://www.kucoin.com/news/flash/polymarket-profitability-crisis-70-of-users-lose-money-profits-concentrated-among-0-04.

This report follows these questions to examine the opportunities and limitations of crypto prediction markets.

The report examines this paradox through three questions. First, how is it possible for a very small number of traders to generate sustained profits on Polymarket? (Section 03) Second, as algorithmic bots rapidly erode excess returns (alpha), does the market transform into a more accurate vehicle for collective intelligence, or does it simply become a domain where only those with information advantages survive? (Section 04) Third, what market design and oversight mechanisms are necessary to preserve the reliability and public-good function of the prices produced by prediction markets? (Section 05)

These three questions are not independent of one another. Taken together, they examine from different angles how the efficiency of price discovery, the distribution of profits, and market integrity interact within a single market structure. Section 06 synthesizes these findings, addresses the core debate, and draws a clear line between fair compensation for improving price discovery and economic rents that require regulatory limits.

The report is primarily intended for policy researchers and regulators examining the introduction and oversight of prediction markets in Korea; journalists and research professionals who cite or use the probability data generated by prediction markets as indicators; and crypto-industry professionals seeking to understand the structure of these markets. It does not examine the business viability of individual platforms or provide investment advice. Its focus is instead on the extent to which, and the conditions under which, the price information generated by prediction markets can be trusted.

Before addressing these questions, Section 02 traces the evolution of prediction markets from informal betting to formalized collective-intelligence markets.

02

From Gambling to Collective Intelligence

To this day, prediction markets have not entirely shed their stigma as little more than a sophisticated form of gambling. What, then, distinguishes prediction markets from traditional betting? The key difference lies in who determines the price and how participants trade. These differences become clear when we divide the ways of staking money on future events into three basic models.

The first is the traditional bookmaker model. The operator sets the odds for each event, takes the other side of participants' bets, and acts as the house. To protect itself from insolvency, the operator builds a conservative margin into the odds, while payout rules and resolution criteria are determined by the house's internal rules.[7] Participants therefore have to place their trust in the operator's financial solvency and integrity.

The second is the centralized exchange-style prediction market. Platforms such as Kalshi and PredictIt use a central server to match buyers and sellers and charge transaction fees. Unlike a bookmaker, however, the operator does not set the odds; participants determine prices through their bids and offers. Transaction records, custody of funds, and settlement, however, still depend on the operator's centralized infrastructure.[8]

The third is the hybrid decentralized prediction market, represented by Polymarket. To enable rapid execution, orders are matched off-chain through a central limit order book (CLOB) hosted on a central server, while records of executed trades and settlement are recorded on-chain on the Polygon blockchain.[9] Event resolution is also handled by a decentralized oracle rather than by the operator. On Polymarket, anyone can propose an outcome after an event concludes by posting a bond. If no one disputes the proposal, the outcome is finalized. If a dispute is repeatedly escalated, final determination is made through a vote by token holders using UMA's (Universal Market Access) Data Verification Mechanism (DVM).[10]

7) Under the bookmaker model, the operator sets the odds and acts as the direct counterparty to participants, with costs incorporated into the odds through the operator's margin. Reichenbach and Walther distinguish bookmakers from exchange-style and decentralized prediction markets by comparing their price-setting mechanisms, counterparties, costs, and settlement structures. F. Reichenbach and M. Walther, "Exploring Decentralized Prediction Markets: Accuracy, Skill, and Bias on Polymarket," SSRN Working Paper No. 5910522 (2025), 7, https://ssrn.com/abstract=5910522.

8) In centralized exchange-style prediction markets, prices are determined by participant orders rather than being set directly by the operator. The platform matches trades and handles fund custody and settlement. Reichenbach and Walther classify PredictIt and Kalshi as belonging to this category. Reichenbach and Walther, "Exploring Decentralized Prediction Markets," 7.

9) Polymarket uses a structure in which orders are matched through a central limit order book (CLOB) and executed trades are settled on the Polygon blockchain. Polymarket Documentation, "Overview," accessed September 5, 2026, https://docs.polymarket.com/trading/overview.

Because settlement is conducted in U.S. dollar-pegged stablecoins (USDC), the binary contract structure—paying $1 if correct and $0 if incorrect—is preserved. This allows the market price to be interpreted as a compressed estimate of the market's view of an event's probability. However, the price should not be interpreted as a pure probability. Transaction costs and institutional constraints—including liquidity, bid-ask spreads, capital locked up until settlement, oracle dispute risk, and differences in regulatory access—are also reflected in the price.[11] In short, decentralized prediction markets build trust through code and collateral rather than the operator's integrity, while stablecoins enable participation across borders.

TABLE 01 **Comparison of Three Types of Future-Event Betting Markets**

Feature	Decentralized Prediction Market	Centralized Exchange-Style	Bookmaker
Examples	Polymarket	Betfair, PredictIt, Kalshi	Bet365, DraftKings
Price-Setting Mechanism	Users or market makers submit orders	Users submit orders	Bookmakers set the odds
Counterparty	Other users (with collateral, token issuance, and settlement via smart contracts)	Other users	Bookmaker (house)
Funds & Collateral	Cryptocurrency (stablecoins)	Fiat currency or cryptocurrency	Fiat currency
Settlement	Oracle reports the event outcome	Exchange rules	Bookmaker rules
Costs	Platform fees and network gas fees	Fees based on net revenue or transaction fees	Included in the odds (overround)
Transparency	Public order book, all on-chain transaction records, and oracle verification	Public order book; transaction data and settlement depend on the operator	Publicly posted odds, opaque margins

Source: Reconstructed by the author based on the market-type classification in Reichenbach & Walther (2025), p. 7. The three-column structure follows the original table, while the definitions and descriptions under "Funds & Collateral" and "Transparency," as well as the examples of platforms, have been supplemented and expanded in this report.

Another key distinction between Kalshi, PredictIt, and Polymarket lies in their regulatory status and participant base. The CFTC designated Kalshi as a Designated Contract Market (DCM) on November 3, 2020. Today, users from most jurisdictions can trade on the platform, with only a limited number of countries and regions restricted. PredictIt initially operated under a 2014 CFTC No-Action Letter granted to Victoria University of Wellington. Although that letter was revoked in 2022, its regulatory basis was restored in modified form through CFTC Letter No. 25-20 in 2025. Polymarket, meanwhile, was ordered in January 2022 to pay a $1.4 million civil monetary penalty and wind down its unregistered event-based binary options markets.[12] As a result, during the 2024 presidential election, Polymarket effectively became the only large-scale global prediction-market platform to bypass geographic restrictions through stablecoins and on-chain infrastructure. Kalshi and PredictIt opened the door to the legitimization of prediction markets under official U.S. regulatory oversight, while Polymarket's 2024 U.S. presidential election market demonstrated collective intelligence in practice.

10) Polymarket uses UMA's (Universal Market Access) Optimistic Oracle to determine market outcomes. Anyone can propose an outcome by posting a bond, typically worth $750. If no dispute is raised during the two-hour challenge period, the proposed outcome is finalized. If the initial proposal is disputed, a new proposal process begins. If the second proposal is also disputed, the matter is referred to UMA's Data Verification Mechanism (DVM), where token holders vote to determine the outcome. Polymarket Documentation, "Resolution," accessed September 5, 2026, https://docs.polymarket.com/concepts/resolution.

11) F. Reichenbach and M. Walther, "Exploring Decentralized Prediction Markets: Accuracy, Skill, and Bias on Polymarket," SSRN Working Paper No. 5910522 (2025), https://ssrn.com/abstract=5910522.

According to an on-chain analysis by Tsang & Yang (2026), Polymarket's 2024 election market exhibited a textbook pattern of market maturation over a ten-month period. Early in the election cycle, liquidity was shallow, meaning that relatively small amounts of capital could move prices significantly. The authors measured this using Kyle's λ (lambda), a price-impact coefficient that indicates how much a given order volume moves the price. The coefficient fell from 0.53 early in the election to 0.01 in the final stages—a fiftyfold decline. Over the same period, the half-life of arbitrage-generating price discrepancies fell from several hours to less than one minute.[13)]

Importantly, this does not mean that arbitrage opportunities disappeared. In fact, a separate study estimated that arbitrageurs earned approximately $40 million in profits on Polymarket over the year beginning in April 2024.[14)] Market maturation does not mean that price discrepancies disappear altogether. Rather, it means that the time during which those discrepancies persist is compressed from hours to mere seconds. Who captured these increasingly fleeting opportunities is the question that Sections 03 and 04 will address.

Deeper liquidity also made the market more resilient to manipulation. In a shallow market, a large trader—or "whale"—can push the price in one direction by deploying substantial capital. In a deep market, however, a distorted price simply creates an opportunity for other participants to take an accurate position at a favorable price. As a result, attempts at manipulation are more likely to end with the manipulator's capital flowing into the market rather than successfully distorting it. This was also the conclusion reached by Tsang & Yang regarding the controversial buying activity of a large account in October 2024, when a French trader named Théo placed massive bets on Trump. They found that capital was flowing simultaneously into both Trump and Harris positions, suggesting that the activity reflected trading between participants with opposing beliefs rather than one-sided market manipulation.[15)]

And the prices produced by this mature market were indeed accurate. Analyzing Polymarket's overall data, Reichenbach & Walther (2025) found that market prices closely tracked realized probabilities across most periods. Compared with online bookmaker odds for the same sporting events, Polymarket prices were also more accurate in 68.2% of cases.[16)]

12) Commodity Futures Trading Commission, "Designated Contract Markets: Kalshi," accessed September 5, 2026, https://www.cftc.gov/IndustryOversight/IndustryFilings/TradingOrganizations/42993; Kalshi, "Can I Trade on Kalshi from Outside the United States?," March 20, 2026, https://help.kalshi.com/en/articles/14026044-can-i-trade-on-kalshi-from-outside-the-united-states; Commodity Futures Trading Commission, "CFTC Letter No. 25-20," July 14, 2025, https://www.cftc.gov/csl/25-20/download; Commodity Futures Trading Commission, "CFTC Orders Event-Based Binary Options Markets Operator to Pay $1.4 Million Penalty," January 3, 2022, https://www.cftc.gov/PressRoom/PressReleases/8478-22.

13) For changes in Kyle's lambda, see p. 16; for changes in the half-life of price deviations, see p. 21. K. P. Tsang and Z. Yang, "The Anatomy of a Blockchain Prediction Market: Polymarket in the 2024 U.S. Presidential Election," arXiv preprint, arXiv:2603.03136 (2026): 16, 21, https://arxiv.org/abs/2603.03136.

14) O. Saguillo et al., "Unravelling the Probabilistic Forest: Arbitrage in Prediction Markets," in *6th Conference on Advances in Financial Technologies (AFT 2025)*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2025): 10, 21, https://doi.org/10.4230/LIPIcs.AFT.2025.27.

15) Tsang and Yang, "The Anatomy of a Blockchain Prediction Market," 21.

16) The comparison of accuracy between Polymarket and online bookmakers is based on the same sample of sports matches. Reichenbach and Walther, "Exploring Decentralized Prediction Markets," 19–21.

However, caution is warranted when generalizing this accuracy across all types of events. Prediction markets perform particularly well in large, high-profile markets with many participants, deep liquidity, and a continuous inflow of public information—such as presidential elections and major sporting events, where attention and capital are concentrated. By contrast, in small markets with limited participation, markets with ambiguous outcome definitions that leave room for settlement disputes, or events in which decisive information is held by only a few participants, prices may not provide reliable probability estimates. In other words, the greater accuracy of prediction markets is an overall tendency, not a characteristic that applies equally to every individual market. Polymarket's consistent projection of a Trump lead during the 2024 U.S. presidential election, when public polls indicated a closely contested race, was precisely a case in which this structural maturity proved its value in a major market where public attention and liquidity were heavily concentrated.

This is precisely what economists had envisioned for prediction markets decades ago. In a 2008 joint paper published in *Science*, twenty-two economists, including Kenneth Arrow, presented prediction markets as a decision-support tool that harnesses collective intelligence.[17] The traditional tools on which governments and corporations rely—opinion polls, expert panels, and bureaucratic committees—are vulnerable to bureaucratic bias, political considerations, and overconfidence among vocal experts. Prediction markets, by contrast, put money at stake, encouraging participants to trade based on the most accurate information available to them rather than on personal reputation or political allegiance. The authors argued that prediction markets could reduce wasteful resource allocation at the national and societal levels and serve as risk-management tools for corporations by allowing them to assess the prospects of new products or projects in advance. At the time, prediction markets were suppressed as a form of private gambling. The authors therefore argued that, given the substantial social value of the information generated by prediction markets, regulators should not suppress them through gambling laws but instead provide safe harbors for low-stakes markets to encourage innovation and the accumulation of data. Sixteen years later, the explosive growth of prediction markets, supported by realworld data, suggests that this vision was not misplaced.

But was the vision of prediction markets as a public good, as envisioned by Arrow and his colleagues, fully realized? The market as a whole clearly became more accurate. Yet the same data reveal another reality: fewer than one-third of traders on Polymarket achieved positive cumulative profits.[18] While the public good of accurate prices was being produced, the profits generated in the process flowed to a small minority. Who are these traders, and what inefficiencies did they identify and exploit to generate their returns?

17) Arrow et al., "The Promise of Prediction Markets," 877–878.

18) Reichenbach and Walther, "Exploring Decentralized Prediction Markets," 20–21.

03

Anatomy of Profit Strategies: Where Do Market Inefficiencies Come From?

In the previous section, we examined the definition of crypto prediction markets and the social value they can provide. However, no market is perfectly efficient from the outset. Technologies mature through repeated cycles of exploitation and refinement, and crypto prediction markets are no different. This section examines where profits actually arise on Polymarket—that is, the sources of market inefficiency. It also uses empirical data to identify the profits generated by gaps in market structure, independently of the ability to make accurate predictions.

The Technical Structure of Prediction Markets: Split, Merge, and Convert

To understand profit strategies, we first need to briefly examine how prediction markets work at the technical level. Markets on Polymarket can broadly be divided into two categories. The first is a Single-Condition Market, in which the outcome is either YES or NO, such as “Will Trump win?” The second is a Multi-Condition Market, in which there are multiple mutually exclusive outcomes, such as “Which team will win the NBA championship: Oklahoma City, San Antonio, or Philadelphia?” Polymarket refers to the latter as a Negative Risk (NegRisk) Market. Regardless of the type, the underlying design principle is the same: the prices of all possible outcomes should add up to exactly $1.00. When the market resolves, only one outcome becomes the winning outcome, and holders of that token receive $1.00 per token.

In a binary market, for example, a YES token for “Trump Wins” priced at $0.60 and a NO token priced at $0.40 add up to exactly $1.00. The same principle applies to a multi-condition market: if there are only four possible outcomes, their prices might be $0.30 + $0.25 + $0.25 + $0.20 = $1.00. In actual markets with many possible outcomes, this total is distributed across numerous contracts. In the NBA market shown in Figure 1, the four teams displayed at the top of the screen—priced at $0.22, $0.20, $0.135, and $0.09—add up to 64.5 cents, while the remaining 35.5 cents is distributed among the other teams shown further down. This, however, is only a design-level identity. In actual order books, the sum frequently deviates slightly above or below $1.00. These deviations are the source of the arbitrage opportunities examined in the next section.

FIGURE 01 **Polymarket NBA 2027 Champion Market Trading Interface**

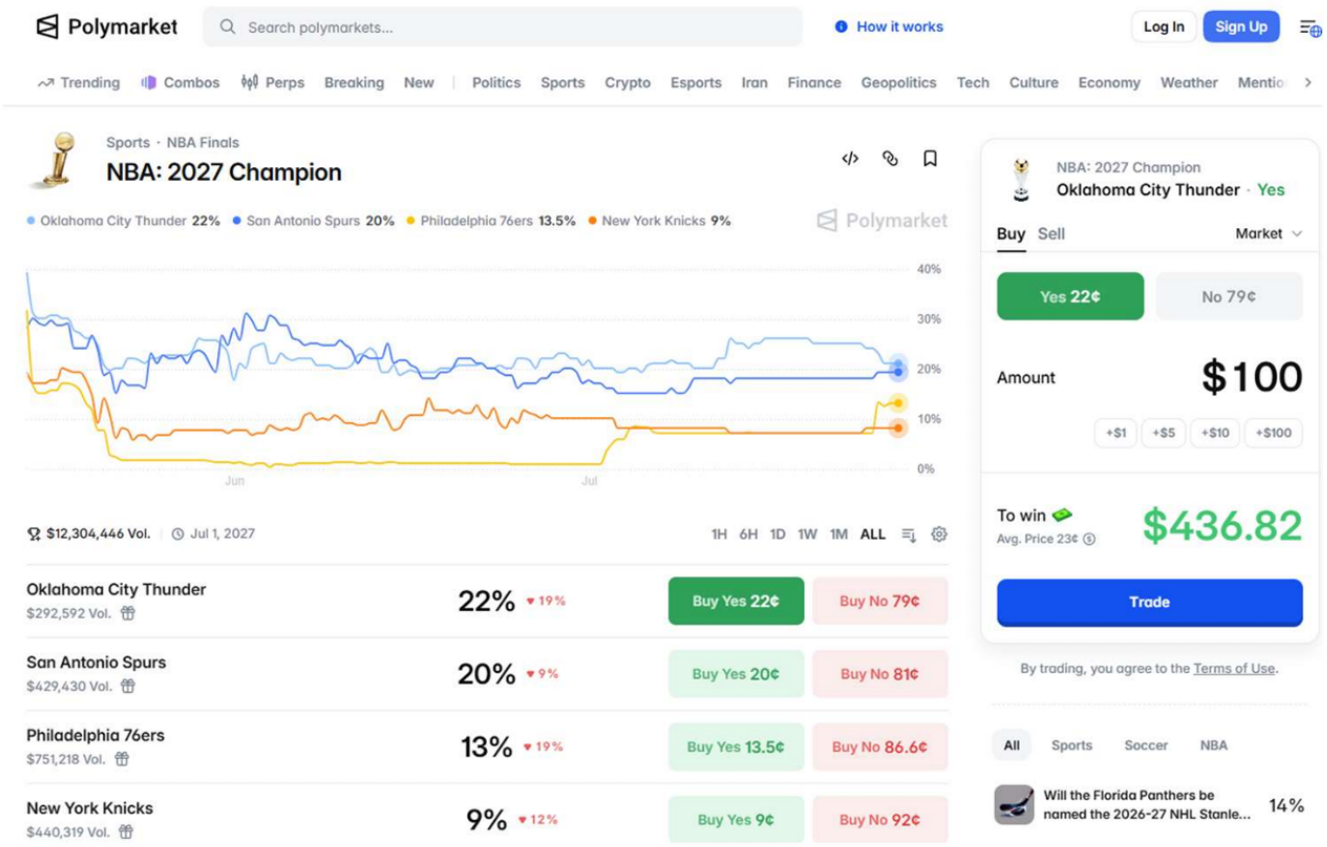

Source: Polymarket (polymarket.com), screenshot taken July 27, 2026. This screen is cited to illustrate the market structure. Rights to the Polymarket trademark and screen design belong to the respective owner.

The trading process for these two types of markets proceeds as follows. ① Users place bets through the web interface, while market makers and arbitrageurs submit orders through an Application Programming Interface (API). API users can also directly call the split, merge, and convert functions as needed. ② These orders are matched off-chain. If someone is already offering to sell the same outcome token, that existing token is transferred directly to the buyer. However, if a user wants to buy YES but no one is offering to sell YES, the order can instead be matched with a user seeking to buy NO, with the two orders executed at a point where their prices add up to $1.00. ③ Because there is no existing token to transfer in this case, the protocol uses the split function to lock $1.00 in collateral and mint new YES and NO tokens, which are then delivered to the respective buyers. This transaction is recorded on the blockchain (on-chain). ④ Traders who want to close their positions before an event is resolved typically place sell orders on the order book. However, when sell orders for YES and NO arrive simultaneously from both sides, the protocol may use the merge function instead of matching the sell orders, burning the two tokens and recovering the locked $1.00 collateral, which is then distributed to the respective sellers. ⑤ When the event is resolved, the oracle (UMA) verifies the outcome. ⑥ Users holding the winning outcome token can redeem it through a smart contract and receive $1.00 per token. Figure 02 summarizes this process, the five ways in which outcome tokens can be processed, and the division of labor between off-chain order matching and on-chain settlement.

During the trading process, outcome tokens can be handled in five ways: Direct Transfer, Split, Merge, Convert, and Unredeemed. A direct transfer means transferring an ERC-1155[19] token as is from one wallet to another, while unredeemed refers to the state in which outcome tokens remain in a user's possession without being redeemed even after the event has been resolved.

FIGURE 02 **Polymarket Trading Lifecycle: The Division of Labor Between Off-Chain Order Matching and On-Chain Settlement**

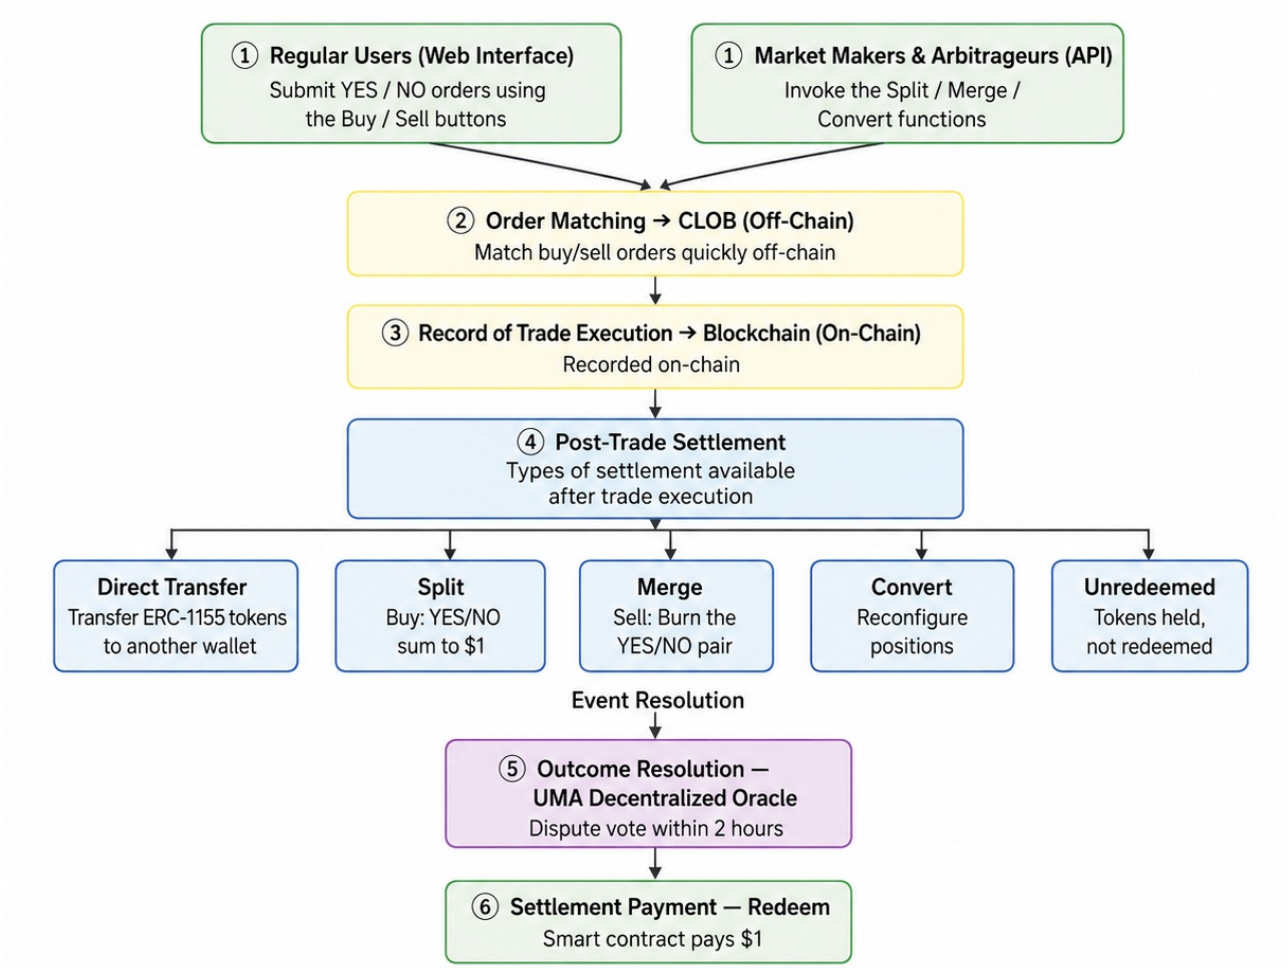

Regular users and API users meet in the same order book, while Split, Merge, and Convert operate at the on-chain settlement layer when needed.

The three core smart contract operations work as follows:

- Split: An operation that locks $1.00 in USDC as collateral and mints a new pair of YES and NO tokens. For example, if a buy order for YES at $0.60 is matched with a buy order for NO at $0.40, bringing the total to $1.00, the protocol locks $1.00 in collateral and issues the two newly minted tokens to their respective buyers.
- Merge: The reverse of a split operation. It burns a pair of YES and NO tokens and releases the $1.00 in locked collateral. This is used when matching sell orders for both sides coincide.

19) ERC-1155 is an Ethereum token standard designed to manage multiple token types within a single smart contract.

- Convert: A function that operates in multi-condition markets, converting a NO position on one candidate into YES positions on all the remaining candidates. Because the sum of the probabilities of all possible outcomes is fixed at 1.00, positions can be converted instantly on-chain without going through buy and sell orders on the order book. For example, a NO position on “Trump Wins” can be converted directly into a set of YES positions on all the other candidates in the market.

The Polymarket interface visible on the web consists simply of Buy and Sell buttons. Arbitrageurs and market makers who trade directly through the API, however, can also use these three functions to execute much more precise and sophisticated orders.

Three Pathways of Logical Arbitrage

The origin of numerous profit strategies lies precisely at this point. The principle that the sum of prices across all possible outcomes should equal exactly $1.00 is only a theoretical ideal; in practice, live order book prices frequently deviate from this equation. This divergence stems from the market’s structural design.

First, the YES and NO tokens associated with a single event, as well as multiple logically interconnected markets, are traded on separate order books. Because the price on each order book is determined solely by the orders submitted to that particular market, price relationships across markets cannot remain in equilibrium on their own unless participants actively correct them through arbitrage or other forms of trading. Second, in markets with shallow liquidity, even a small number of orders can move prices significantly, causing large price discrepancies to arise frequently. Third, ordinary participants tend to focus on the individual market they are trading rather than monitoring the logical relationships across the market as a whole in real time. The role of arbitrageurs is to identify these gaps and bring the prices back into alignment. The resulting price differential is their profit.

These opportunities are often referred to as risk-free arbitrage, but this is a theoretical term. In practice, traders face various risks and costs, including leg risk, where one side of a trade is executed while the other side is delayed or the price moves; the opportunity cost of capital locked up until settlement; the possibility of an oracle dispute; and restrictions on access to the platform. Taking these factors into account, this report refers to such strategies as theoretically near-risk-free logical arbitrage, or simply logical arbitrage.

Saguillo et al. (2025) classify arbitrage within Polymarket into two types: rebalancing arbitrage and combinatorial arbitrage. Drawing on Ng et al. (2026)’s study comparing prices across platforms, this report adds cross-platform arbitrage as a third pathway.

The first type is single-condition rebalancing arbitrage. It occurs when the combined prices of YES and NO within a single condition deviate from $1.00. For example, if “Trump Wins” YES is trading at $0.80 and NO at $0.10, an arbitrageur can buy both for a combined $0.90 and receive $1.00 when the market resolves, locking in a guaranteed profit of $0.10. Conversely, if the combined price exceeds $1.00, the arbitrageur can use the split function to mint a YES/NO token pair for $1.00 and immediately sell both tokens to capture the difference.

Saguillo et al. analyzed all 17,218 individual Polymarket events that resolved between April 2024 and April 2025. They found at least one instance of this type of opportunity in 7,051 events, or approximately 41% of the total. The median theoretical return per $1.00 of capital allocated to an opportunity was approximately 60 cents.[20] However, this figure should not be interpreted directly as a rate of return without taking three important qualifications into account.

First, this is a median based on the number of opportunities, not on the amount of capital involved. In other words, it does not mean that an average $1.00 invested earned $0.60. Rather, if all the opportunities observed in the order books were ranked by theoretical return, the opportunity at the midpoint of that distribution offered a theoretical return of 60%.

Second, these extreme price discrepancies were not evenly distributed across the market but tended to cluster at specific points. These included contracts whose prices had fallen to a few cents because the outcome was effectively decided, periods immediately before market resolution when order book liquidity disappeared, and small markets with little trading activity. In such situations, theoretical returns could reach several hundred percent, but the amount of capital that could actually be deployed was small because the order books were thin.

Third, high percentage returns and a small aggregate profit can coexist without contradiction. In fact, the total extracted profit from all such opportunities over the one-year period was approximately $39.6 million, as discussed in the following section, “Who Took How Much?”—a relatively small amount compared with the platform’s total trading volume over the same period.

Nevertheless, this figure is worth highlighting not because of its dollar value, but because of what it reveals about the state of the market. In mature equity and futures markets, a theoretical arbitrage opportunity offering a 60% return is extremely rare, and even when such an opportunity appears, it typically disappears within milliseconds. On Polymarket, by contrast, such opportunities were observed at least once in 41% of all markets. In short, the figure does not mean that “anyone could have earned 60%.” Rather, it indicates how frequently price relationships were left out of alignment. Whether an opportunity could actually be exploited depended heavily on the available order volume, how long the opportunity lasted, execution slippage, and the period for which capital remained locked up. In an extreme example, one user (@Tutaaa91) simultaneously accumulated YES and NO tokens for less than $0.02 each, locking in a guaranteed profit of $58,983 from a single trade.[21]

The second type is combinatorial arbitrage, which occurs when the combined YES prices across all outcomes in a multi-condition market deviate from $1.00, or when the prices of two logically dependent markets contradict one another.

For example, consider two markets asking “Which party will win?” and “By how many percentage points will it win?” If “Republicans Win by 2 percentage points or more” resolves YES, then “Republicans Win” must also resolve YES. This establishes a logical lower and upper bound between the prices of the two contracts. If this relationship breaks down, a trader can take opposing positions across the two markets and construct a portfolio that wins regardless of the outcome.

20) The median is calculated across the observed arbitrage opportunities rather than as a return weighted by the amount invested in each opportunity. Saguillo et al., “Unravelling the Probabilistic Forest,” 10.

21) Saguillo et al., “Unravelling the Probabilistic Forest,” 16–18.

Saguillo et al. identified intra-market rebalancing opportunities in 662 of 1,578 multi-condition markets, with approximately 100 such opportunities per market on average.[22] In markets related to the 2024 U.S. presidential election, they also used Large Language Model (LLM)-based semantic analysis to identify 13 pairs of logically dependent markets and confirmed that actual arbitrage was executed in 5 of those pairs.[23]

The third type is cross-platform arbitrage. This strategy targets price discrepancies that arise when contracts on the same event are traded simultaneously across multiple platforms, such as Polymarket, Kalshi, and PredictIt, because of differences in regulation and participant composition.

According to Ng et al., during the 2024 presidential election, the Harris victory contract traded on Kalshi at an average of 3.27 cents more than on Polymarket, while the Trump victory contract traded on Polymarket at an average of 2.59 cents more than on Kalshi. These price differences remained almost continuously in the same direction throughout the 33-day period during which the two platforms overlapped.[24]

When trade data from the two platforms were matched at one-second intervals, the researchers found approximately $130,000 in arbitrage opportunities during the final 15 days before the election, compared with up to $2.2 million using a five-minute matching interval.[25] During the 2024 election, the two platforms differed in the regions where they were accessible and in their regulatory conditions. This market segmentation allowed price differences to persist for contracts tied to the same event. In other words, the Law of One Price did not fully hold across the two platforms.

Who Took How Much?

In their empirical study of two types of arbitrage within Polymarket, Saguillo et al. (2025) estimate that approximately $39.6 million in profits were extracted over the one-year observation period. This estimate does not include profits from cross-platform arbitrage.[26] The distribution of these profits was also extremely skewed. The single top-performing account captured approximately $2.01 million through 4,049 orders, while each of the top ten accounts earned profits ranging from several hundred thousand dollars to more than $2 million. Some of the top accounts participated in thousands of orders, leading the researchers to characterize this high-frequency activity as “bot-like behavior.”[27]

22) Saguillo et al., “Unravelling the Probabilistic Forest,” 11–12.

23) Saguillo et al., “Unravelling the Probabilistic Forest,” 14–15, 21.

24) Throughout the 33-day overlapping trading period, the Kalshi price for the Kamala Harris victory contract consistently exceeded the corresponding price on Polymarket. H. Ng, L. Peng, Y. Tao, and D. Zhou, “Price Discovery and Trading in Modern Prediction Markets,” SSRN Working Paper No. 5331995 (2026), Table 3, https://doi.org/10.2139/ssrn.5331995.

25) The aggregation period spans from October 23 to November 5, 2024. Before transaction costs, total arbitrage profits amounted to approximately $130,000 under one-second matching and up to $2.2 million under five-minute matching. Assuming transaction costs of 2%, net arbitrage profits fell to approximately $41,000 and $850,000, respectively. Ng et al., “Price Discovery and Trading in Modern Prediction Markets,” Table 4.

26) The precise estimate of arbitrage profits is $39,587,585, based on a conservative aggregation that includes only trades generating at least $1 in profit. This figure covers only profits extracted through arbitrage strategies and therefore measures a different population from the overall net profits and losses of all traders. Saguillo et al., “Unravelling the Probabilistic Forest,” § 7.4.

27) The estimated profit of the top account was $2,009,631.76, and the dataset analyzed comprised a total of 86,620,143 executed orders. Saguillo et al., “Unravelling the Probabilistic Forest,” § 7.1 and Table 1.

The picture is much the same when we broaden the scope beyond arbitrage. An on-chain analysis that aggregated the net profits and losses of approximately 1.73 million Polymarket wallet addresses found that fewer than the top 0.04% of addresses captured more than 70% of total net profits.[28)]

The two figures in this section, however, should be distinguished because they come from different sources and measure different things. The $39.6 million figure was calculated by Saguillo et al. by identifying and aggregating only those transactions in Polymarket's execution data that qualified as arbitrage. It therefore measures "the profits generated by a specific strategy." By contrast, the 0.04% and 70% figures aggregate net profits and losses across all strategies at the wallet-address level, measuring "the overall financial performance of all participants." These are not two slices of the same pie; they are measurements of two different pies.

Nevertheless, it is worth considering them together because, whether viewed at the level of individual strategies or at the level of participants, the evidence points in the same direction. Arbitrage profits were concentrated among a small number of top accounts, while overall net profits were likewise disproportionately concentrated among a small group of top wallet addresses.

Section Summary: Inefficiencies That Will Fade With Market Maturity, and the Questions That Remain

In sum, the primary arbitrage profits examined in this section stem from structural gaps created by insufficient liquidity, wide bid-ask spreads, and cross-platform fragmentation. These types of inefficiencies are likely to diminish substantially as liquidity deepens and the market matures. Indeed, the time windows during which arbitrage opportunities remain open are already shrinking rapidly.

This raises the question: where will persistent profit advantages come from once these structural inefficiencies have been reduced? The challenge is that even after pricing gaps close, other forms of information advantage—such as insider information, outcome manipulation, and differences in the interpretation of settlement rules—may remain.

28) This aggregation covers total net profits and losses at the individual wallet-address level and therefore differs from the arbitrage profit estimate in Footnote 26 in both what is measured and the period covered. KuCoin, "Polymarket Profitability Crisis."

04

The Decay of Alpha and the Transformation of Collective Intelligence

The logical arbitrage examined in Section 03 represented a primary source of structural profits observed on Polymarket.

However, those market gaps are now closing rapidly. This section first examines the process of "alpha" decay through empirical data, then explores the issues that remain even after these gaps have been fully closed—namely, insider information, outcome manipulation, and the influence of settlement rules. Finally, it examines how closely this transition resembles the path that traditional equity markets traversed first.

The Decay of Alpha: Closing Profit Windows

The maturation of prediction markets does not merely improve price accuracy; it also changes the rate at which profit opportunities arising from price discrepancies disappear. Analyzing the complete on-chain dataset from the 2024 U.S. presidential election markets, Tsang and Yang (2026) empirically demonstrated this trajectory. In early 2024, when the combined price of YES and NO tokens deviated from $1.00, it took several hours for the discrepancy to be reduced by half. By October–November 2024, however, as market liquidity deepened, this half-life had fallen to less than one minute.[29] Price discrepancies remained a source of arbitrage profits, but the time available to capture and monetize them was shrinking just as rapidly.

Kyle's λ (lambda), a price-impact coefficient that measures how much a given order volume moves the price, also fell sharply as the market deepened, declining from 0.53 to 0.01—roughly one-fiftieth of its previous level.[30] Lambda is fundamentally a market microstructure measure of illiquidity, but because it reflects the extent to which an individual participant's order can affect the market price, it can also be interpreted as a proxy for vulnerability to manipulation.

Arbitrageurs and market makers were at the center of this transformation. Using automated bots, they identified and closed price discrepancies faster than other participants. As competition among bots to capture these opportunities intensified, the scope for earning excess returns rapidly diminished.

29) Tsang and Yang, "The Anatomy of a Blockchain Prediction Market," 14–16.

30) Tsang and Yang, "The Anatomy of a Blockchain Prediction Market," 16.

Yet while the size of these returns declined, the risks and costs associated with trading remained. Although Polymarket did not charge per-trade fees during the period analyzed,[31] arbitrage was by no means costless. Because of the non-atomicity of order execution, one side of a trade could be filled while the price on the other side moved adversely before it was executed. Traders also incurred opportunity costs because capital remained locked up until trades were completed or positions were settled.

As competition among bots reduced expected margins, these costs increasingly eroded net arbitrage profits. Eventually, simply being the first to identify a market gap was no longer enough to generate excess returns.

This, however, does not mean that trading skill has disappeared. Analyzing more than 124 million trades on Polymarket, Reichenbach and Walther found that market prices generally tracked realized probabilities with a high degree of accuracy, and were slightly more accurate than forecasts based on traditional bookmaker odds. Within this overall market efficiency, however, the profits of some traders exhibited persistence that could not be explained by luck alone. Their returns reflected an ability to exploit the systematic biases of less sophisticated participants, including over-trading of default options and a tendency to favor YES positions.[32]

In short, as the alpha from mechanical arbitrage rapidly declines, the ability to identify and trade against the systematic biases of other participants is emerging as a new source of competitive advantage.

What Remains: Three Types of Information Advantage

What competitive advantages remain after market microstructure inefficiencies have been reduced? They can be divided into three types: those who know the outcome before others do, those who can directly influence the outcome, and those who can influence the definition of the outcome and the settlement rules.

The first type is the classic insider—someone who knows the outcome of an event before others do. Prediction markets are more broadly exposed to the problem of insider information than traditional financial markets because they offer contracts on events such as wars, policy decisions, and military operations, where the outcome may be known in advance by only a small number of people.

These concerns have already materialized. In February 2026, Israeli prosecutors indicted a reserve military officer and a civilian accomplice. Investigators found that the reserve Air Force major had obtained classified information about the timing of Israel's strikes on Iran—Operation Rising Lion, also known as the "12-Day War" of June 2025—during a security briefing two days before the operation and passed it to the accomplice. The accomplice used this information to place bets on Polymarket, earning more than $160,000 in profits, which the two subsequently shared.[33]

31) During the observation period analyzed in this study, Polymarket did not charge per-trade transaction fees. Saguillo et al., "Unravelling the Probabilistic Forest," § 7.1.

32) An analysis of more than 124 million Polymarket trades found that market prices tracked realized probabilities relatively closely and were slightly more accurate than forecasts based on bookmaker odds. In addition, while no clear general longshot bias was observed, biases toward overtrading the default option and the YES option were identified, and skilled traders appeared to profit by exploiting the biases of less-skilled participants. Reichenbach and Walther, "Exploring Decentralized Prediction Markets."

On-chain analytics firm Bubblemaps identified a pattern in which nine interconnected Polymarket accounts achieved a 98% win rate across more than 80 military-related bets—including bets on the date of the U.S. military strikes on Iran, the ouster of Iran's Supreme Leader, and the timing of a ceasefire—and earned more than $2.4 million.[34] In the United States, an active-duty Army soldier was also indicted for allegedly using classified information to place bets on the capture of Venezuelan President Nicolás Maduro.[35]

The second type is a distinct problem from insider information: manipulators who go beyond knowing an outcome in advance and can actually create or influence the outcome themselves. In July 2026, Spotify discovered that more than 500,000 artificial streams had been used to push Malcolm Todd's song "Earrings" toward the top of the U.S. charts. The manipulated ranking had already been used as the settlement criterion for Kalshi's "Most Streamed Song in June" market, in which more than $3 million had been traded. Spotify subsequently demanded that both Kalshi and Polymarket remove its corporate logo and explicitly state that they had no official partnership with Spotify.[36]

This case demonstrates that the integrity of prediction markets depends not only on how effectively they aggregate information, but also on how independent the metrics underlying those markets are from artificial intervention by market participants.

The third type consists of those who can influence the definition of an outcome and the settlement rules. Prediction-market contracts define in words what occurrence constitutes a "YES" outcome and delegate the determination of that outcome to an oracle. When an event unfolds ambiguously, the room for interpretation in that definition and the oracle's dispute-resolution procedures can themselves become sources of advantage.

33) The indictments were made public in February 2026. The fact that the defendant was a reserve major in the Israeli Air Force, along with specific details of the case, became known in March of the same year after a gag order was partially lifted. The account involved generated approximately $163,000 in profits, and the operation in question was Operation Rising Lion, which began on June 13, 2025. "Two Indicted for Using Classified Info to Place Online Bets on Military Operations," *The Times of Israel*, February 12, 2026, https://www.timesofisrael.com/two-indicted-for-using-classified-info-to-place-online-bets-on-military-operations/; Emanuel Fabian, "Israeli Air Force Major Charged with Using Classified Info to Place Bets on Polymarket," *The Times of Israel*, March 26, 2026, https://www.timesofisrael.com/israeli-air-force-major-charged-with-using-classified-info-to-place-bets-on-polymarket/.

34) On-chain analytics firm Bubblemaps analyzed the trading patterns of interconnected accounts across military-related prediction markets. A. Chasan, "Suspected Insider Accounts Net $2.4 Million on Polymarket Iran War Bets with 98% Win Rate, Firm Finds," *CBS News: 60 Minutes*, May 17, 2026, https://www.cbsnews.com/news/betting-on-iran-war-insider-trading-concerns-prediction-markets-60-minutes/.

35) Gannon Ken Van Dyke, a U.S. Army soldier, was indicted on April 23, 2026, for allegedly using classified information related to the operation to capture Nicolás Maduro to generate approximately $410,000 in profits through prediction market trades. U.S. Department of Justice, "U.S. Soldier Charged with Using Classified Information to Profit from Prediction Market Bets," April 23, 2026, https://www.justice.gov/opa/pr/us-soldier-charged-using-classified-information-profit-prediction-market-bets; J. Yurcaba, "U.S. Soldier Charged with Insider Trading over Maduro's Ouster," *NPR*, April 23, 2026, https://www.npr.org/2026/04/23/nx-s1-5797957/maduro-raid-charges-polymarket-insider.

36) Spotify announced that it had removed more than 500,000 artificial streams and would introduce additional verification procedures before publishing its chart rankings. As of the time of reporting, there was no evidence that the artist or the artist's team had been directly involved in the manipulation. D. Tsekova, "Spotify Challenges Prediction Markets After Song Chart Rigging," *Bloomberg*, July 2, 2026, https://www.bloomberg.com/news/articles/2026-07-02/spotify-challenges-prediction-markets-after-song-chart-rigging.

Polymarket's UMA oracle finalizes outcomes through bond-backed proposals and token-holder voting. In markets where the outcome is disputed, participants who understand the rules of the dispute process and the distribution of token voting power are structurally advantaged. While the first two types derive their advantage from the real-world events underlying the market, the third derives its advantage from the internal rules governing contracts and settlement.

All three, therefore, involve problems at a layer that speed competition, as discussed in the first subsection of Section 04, can never fully resolve.

On-Chain Transparency: Insiders Are Surprisingly Visible

The fact that prediction markets are vulnerable to insider information does not mean that the trading footprints of insiders are difficult to detect.

In some respects, tracking insiders is actually easier in prediction markets than in traditional markets, for two reasons.

First, all executed trades on Polymarket are recorded as public records on the blockchain. Although the identities of wallets remain anonymous, anyone can query when a particular wallet traded, which market it traded in, and how much it wagered. This transparency is what enabled private analytics firms such as Bubblemaps to identify clusters of suspicious accounts before government investigations did.[37)]

Second, insiders reveal themselves statistically. High-frequency traders without an information advantage typically achieve win rates only slightly above 50%, whereas the accounts implicated in recent cases posted win rates of 80% to 98%. Repeatedly placing large bets on events with low market prices—that is, high potential payouts—and getting the outcome right cannot easily be dismissed as mere chance. The pattern itself becomes a clue for investigation.[38)]

Academic research points in the same direction. Ng et al. (2026) found that net buying imbalances in large orders significantly predicted subsequent 30-minute returns, showing that the intensity of informed trading can be measured through traces left in price data.[39)] If a particular wallet repeatedly concentrates its bets in one direction immediately before an event occurs, and the market subsequently moves in that direction, such a pattern can provide an important clue that the account possesses an information advantage. Ultimately, the key to effective surveillance is not the technology to detect these traces, but the institutional framework needed to connect them to investigation and enforcement.

37) Chasan, "Suspected Insider Accounts Net $2.4 Million on Polymarket Iran War Bets."

38) Professor Todd Phillips of Georgia State University noted that while typical high-frequency traders achieve win rates only slightly above 50%, sustained win rates in the 80% to 90% range are highly unusual. "Trader Made Nearly $1 Million on Polymarket with Remarkably Accurate Iran Bets," *CNN Politics*, March 24, 2026, https://www.cnn.com/2026/03/24/politics/iran-war-bets-prediction-markets.

39) A one-standard-deviation increase in the net buying imbalance of orders in the top 1% by size on a given trading day was associated with an approximately 20–33 basis-point increase in subsequent 30-minute returns, whereas small orders showed no comparable predictive power. Ng et al., "Price Discovery and Trading in Modern Prediction Markets," Tables 8–9.

The Precedent of Stock Markets

This entire transition—the erosion of alpha through speed competition and the institutionalization of insider-trading surveillance—follows a path that stock markets traversed first. During the heyday of HighFrequency Trading (HFT), approximately $300 million was invested in constructing ultra-low-latency fiber-optic cables connecting Chicago and New York, while trading firms competed in an arms race for speed by using colocation—placing their servers as close as possible to exchange matching engines.[40] However, as competitors achieved similar speeds, the excess returns from speed alone diminished while the costs of the arms race increased. The competition among bots in prediction markets is following a similar trajectory.

Nevertheless, the experience of stock markets should not be viewed as a simple solution to these problems. Stock markets did not eliminate insider trading or high-frequency trading; rather, they have managed these issues institutionally by establishing rules and supervisory systems to define and monitor them. Prediction markets are only at the beginning of this process.

In the United States, lawmakers have introduced the *Banning Event Trading on Sensitive Operations and Federal Functions Act* (BETS OFF Act), which would prohibit event contracts involving government actions, war, and terrorism, as well as the use of non-public information by public officials.[41] Polymarket and Kalshi are now moving toward codifying prohibitions on insider trading and strengthening self-regulatory mechanisms for referring suspicious trading activity to law enforcement agencies.

Furthermore, it would be a mistake to assume that prediction markets will follow exactly the same path as stock markets. Prediction markets trade not corporate ownership stakes, but the events themselves—such as wars, elections, sports, and cultural developments. Because trading can intersect with a broad range of stakeholders directly involved in these events, prediction markets inevitably raise a wider range of ethical and political issues than stock markets.

Before considering insider-trading regulation, we must first ask a more fundamental question: should the outbreak of a war involving human lives be a target for betting at all? This is ultimately a question about the legitimacy of creating such a market in the first place. In short, many of the problems facing prediction markets have already been addressed in stock markets through market maturation and institutional development. Others, however, are unique to prediction markets and require new answers. In addition, the public availability of on-chain transaction records makes it easier to track suspicious trading flows and recurring patterns.

40) Spread Networks' 825-mile ultra-low-latency fiber-optic cable connecting Chicago and New York was estimated to cost approximately $300 million to construct and reduced transmission time by about three milliseconds compared with the previously fastest route. The U.S. Securities and Exchange Commission (SEC) explains that high-frequency trading firms use colocation services, placing their servers as close as possible to exchange matching engines to increase trading speed. Christopher Steiner, "Wall Street's Speed War," *Forbes*, September 9, 2010, https://www.forbes.com/forbes/2010/0927/outfront-netscape-jim-barksdale-daniel-spivey-wall-street-speed-war.html; Mary L. Schapiro, "Testimony Concerning the Severe Market Disruption on May 6, 2010," U.S. Securities and Exchange Commission, May 11, 2010, https://www.sec.gov/news/testimony/2010/ts051110mls.htm.

41) Office of Senator Chris Murphy, "Murphy, Casar Introduce Bicameral Bill to Ban Prediction Markets on Government Actions, War, and Events Ripe for Rigging," March 17, 2026, https://www.murphy.senate.gov/newsroom/press-releases/murphy-casar-introduce-bicameral-bill-to-ban-prediction-markets-on-government-actions-war-and-events-ripe-for-rigging.

05

Institutional Design to Safeguard Prediction Market Integrity

The public-good functions of prediction markets are not guaranteed by market growth alone. If participants can directly manipulate the underlying metrics of a market, or if those with privileged information trade before that information is reflected in prices, market prices become records of manipulation and information leakage rather than genuine forecasts. To preserve the credibility of the probability signals generated by prediction markets, their integrity must therefore be protected through appropriate market design and oversight. What institutional safeguards are needed to achieve this?

First, prediction markets need clear prohibitions against insider trading and effective surveillance systems. In April 2026, Polymarket partnered with Chainalysis to introduce a transaction-monitoring system. The previous month, it revised its Market Integrity Rules to prohibit trading based on stolen confidential information, unlawful insider information, or the ability to influence an outcome.[42] Kalshi likewise already enforces insider trading rules comparable to those applied by the New York Stock Exchange (NYSE) and Nasdaq.[43]

However, platform self-regulation alone is not sufficient. As discussed in Section 04, public on-chain records can make suspicious trading patterns particularly visible. The key issue is establishing procedures that can connect detected anomalies to investigation and sanctions. The scope of prohibited use of insider information, platforms' reporting and cooperation obligations, and liability for violations should therefore be clearly defined.

Second, platforms need stronger listing reviews for metrics that can be manipulated. As the Spotify case discussed in Section 04 illustrates, if market participants can directly purchase or manipulate the metric on which a contract is based, the market ceases to function as a prediction mechanism and instead becomes a self-fulfilling order.

The manipulability of the underlying metric, the cost of manipulation relative to the amount wagered on the market, and the objectivity of outcome verification should all be key criteria in listing reviews. Platforms should also be able to refuse to list categories of markets with a high risk of manipulation.

42) Polymarket published Market Integrity Rules in March 2026 addressing trading based on insider information and the ability to influence market outcomes, and in April partnered with Chainalysis to deploy an on-chain market integrity monitoring system. Polymarket, “Polymarket Publishes Enhanced Market Integrity Rules Across Its DeFi Platform and CFTC-Regulated U.S. Exchange,” March 23, 2026, https://www.businesswire.com/news/home/20260320997513/en/Polymarket-Publishes-Enhanced-Market-Integrity-Rules-Across-Its-DeFi-Platform-and-CFTC-Regulated-U.S.-Exchange; Polymarket, “Polymarket Selects Chainalysis to Deploy First-of-Its-Kind On-Chain Market Integrity Solution,” April 30, 2026, https://www.businesswire.com/news/home/20260430726176/en/Polymarket-Selects-Chainalysis-to-Deploy-First-of-Its-Kind-On-Chain-Market-Integrity-Solution.

43) Kalshi prohibits trading by individuals in possession of material non-public information (MNPI), and its CEO has stated that the rule was modeled on the insider trading rules of the New York Stock Exchange (NYSE) and Nasdaq. B. Metzger, “Kalshi CEO Says Insider Trading Is ‘Committing a Financial Crime’ as He Backs New Bill,” *Yahoo Finance*, January 7, 2026, https://finance.yahoo.com/news/kalshi-backs-bill-bar-prediction-165252701.html.

Third, oracle dispute-resolution procedures must be made more transparent and fair. Under Polymarket's UMA oracle system, anyone can propose an outcome by posting a bond typically worth about $750. If no objection is raised within two hours, the outcome is finalized. If an objection is raised, a new proposal process begins. If the second proposal is also disputed, the case is escalated to UMA's Data Verification Mechanism (DVM), where token holders vote to determine the outcome.[44)]

However, this process raises the possibility that token holders' financial interests in a particular outcome could influence the decision. Safeguards are therefore needed to prevent conflicts of interest among those involved in the resolution process, along with transparent deliberation procedures and mechanisms for reviewing decisions after they have been made.

Fourth, restrictions are needed on trading by public officials and on the listing of sensitive events. As illustrated by the Israeli Air Force major's case discussed in Section 04, betting based on classified information obtained through one's official duties fundamentally undermines market integrity.

Rules could require public officials to report and disclose their prediction market transactions or prohibit them altogether from trading in markets related to their official responsibilities. At the same time, clear criteria are needed to exclude events directly affecting social trust or national security—such as the timing of the outbreak of a war or the likelihood of an assassination—from being listed as prediction market contracts.

44) UMA’s Data Verification Mechanism uses economic guarantees and a dispute resolution process to determine outcomes. On Polymarket, if a dispute arises during the designated challenge period following a proposed outcome, the matter proceeds to UMA’s verification process. H. Lambur, A. Lu, and R. Cai, “UMA Data Verification Mechanism: Adding Economic Guarantees to Blockchain Oracles,” Risk Labs Technical Report (April 2020), https://github.com/UMAprotocol/whitepaper/blob/master/UMA-DVM-oracle-whitepaper.pdf.

06

Conclusion: The Paradox of Collective Intelligence

Crypto prediction markets aggregate dispersed information into prices, producing signals that can closely track real-world outcomes. As participation and liquidity have expanded, price discrepancies have been corrected more rapidly, allowing markets to operate more efficiently. However, greater efficiency in price discovery has not made the distribution of profits more equal. Arbitrage profits have been concentrated among a small number of top accounts, and overall net profits have likewise been concentrated among a small number of addresses. Market efficiency cannot, by itself, be regarded as the cause of this profit concentration. Nevertheless, it is clear that efficient price formation and an unequal distribution of profits can coexist.

Grossman and Stiglitz (1980) showed that a perfectly efficient market cannot exist as long as gathering information is costly.[45] If prices already reflect all available information, no one has an incentive to bear the cost of collecting it; but if no one collects information, prices cannot reflect it. For markets to incorporate information into prices, participants who gather information or correct price discrepancies must be able to earn returns that compensate for the costs and risks involved. From this perspective, I view some of Polymarket's arbitrage profits as compensation for the costs and risks involved in identifying and eliminating price discrepancies. The fact that profits are concentrated among a small number of participants does not, by itself, mean that the market has failed.

The question is where to draw the line between legitimate compensation and something else. Competition to invest in faster servers, develop sophisticated bots, and identify price gaps before other participants is, in principle, open to everyone. As competition has intensified, we have also seen that this form of alpha tends to decline. However, privileged access to confidential information, the ability to directly influence the metric on which a market is based, and the power to influence the interpretation or settlement of market rules are fundamentally different. I do not believe that profits derived from such advantages can be regarded as legitimate compensation for price discovery. They are closer to rents arising from the structure of the market than to rewards the market must pay to become more accurate. The four institutional measures proposed in Section 05 are intended to define and manage this boundary.

Let us return to the question posed in Section 01: Is this structure of hyper-concentrated returns a market failure, or is it the natural outcome of information asymmetry?

45) S. J. Grossman and J. E. Stiglitz, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review* 70, no. 3 (1980): 393–408.

The answer offered by this report is that it is the natural result of information advantage. However, an important boundary must be drawn. When participants identify price gaps before others and close them, the greater profits they earn are part of the reward generated by the market's process of arriving at more accurate prices. Arbitrageurs and bots do not cause market failure simply by closing price gaps. On the contrary, their activity helps prices move more quickly toward the actual outcome and improves the quality of the probability signals available to outside observers, including journalists, policymakers, and researchers.

However, not every form of information advantage can be justified in the same way. Profits derived from privileged access to confidential information, the ability to directly influence an outcome, or the authority to intervene in the settlement process are fundamentally different from rewards earned through price discovery. They are not compensation the market must pay to function efficiently, but rents that should be constrained through institutional safeguards. The problem prediction markets must solve, therefore, is not profit concentration itself, but how to distinguish legitimate information advantages from rent-seeking.

The mainstream adoption of prediction markets has already begun. Meta is developing its own prediction market app, Arena; exchanges and securities firms are incorporating prediction contracts into their product offerings; and institutional capital is entering the market. The question is no longer whether prediction markets will become mainstream, but what form that mainstream adoption will take. If comprehensive on-chain records, statistical detection of anomalies, and institutions capable of connecting those findings to investigation and sanctions work together, prediction market data could become more trustworthy and develop into a form of social infrastructure that complements opinion polls and expert forecasts. Conversely, if oversight and market design fail to keep pace with market expansion, we may find ourselves with increasingly sophisticated forms of information asymmetry hidden behind prices that merely appear more accurate.

The paradox of collective intelligence does not mean that prediction markets have failed. Rather, it shows that collective intelligence is not some mystical quality that emerges automatically. Only when legitimate rewards for information advantage are distinguished from rent-seeking, and market integrity is protected from manipulation and insider information, can market prices remain a form of collective wisdom that society can trust.

The challenge this new market leaves us with is to ensure that prediction market prices become collective wisdom rather than a hunting ground for a privileged few.

References

Arrow, K. J., R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. D. Nelson, G. R. Neumann, M. Ottaviani, T. C. Schelling, R. J. Shiller, V. L. Smith, E. Snowberg, C. R. Sunstein, P. C. Tetlock, P. E. Tetlock, and E. Zitzewitz. 2008. "The Promise of Prediction Markets." *Science* 320, no. 5878: 877–78. https://doi.org/10.1126/science.1157679.

Chasan, A. 2026. "Suspected Insider Accounts Net $2.4 Million on Polymarket Iran War Bets with 98% Win Rate, Firm Finds." *CBS News: 60 Minutes*, May 17, 2026. https://www.cbsnews.com/news/betting-on-iran-war-insider-trading-concerns-prediction-markets-60-minutes/.

Coinbase. 2025. "System Update: The Future of Finance Is on Coinbase." December 17, 2025. https://www.coinbase.com/blog/system-update-the-future-of-finance-is-on-coinbase.

CoinDesk. 2026. "Coinbase (COIN) and Kalshi Team for Prediction Market for Crypto Exchange's U.S. Customers." January 27, 2026. https://www.coindesk.com/markets/2026/01/27/coinbase-rolls-out-prediction-market-to-u-s-customers.

Commodity Futures Trading Commission. 2020. "CFTC Designates KalshiEX LLC as a Contract Market." November 4, 2020. https://www.cftc.gov/PressRoom/PressReleases/8302-20.

Commodity Futures Trading Commission. 2022. "CFTC Orders Event-Based Binary Options Markets Operator to Pay $1.4 Million Penalty." January 3, 2022. https://www.cftc.gov/PressRoom/PressReleases/8478-22.

Commodity Futures Trading Commission. 2025. "CFTC Letter No. 25-20: Amendment of CFTC Letter No. 14-130 Regarding the Operation of a Small-Scale, Not-for-Profit Market for the Trading of Event Contracts for Educational Purposes." July 14, 2025. https://www.cftc.gov/csl/25-20/download.

Fabian, E. 2026. "Israeli Air Force Major Charged with Using Classified Info to Place Bets on Polymarket." *The Times of Israel*, March 26, 2026. https://www.timesofisrael.com/israeli-air-force-major-charged-with-using-classified-info-to-place-bets-on-polymarket/.

Grossman, S. J., and J. E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70, no. 3: 393–408.

Kalshi. 2026. "Can I Trade on Kalshi from Outside the United States?" *Kalshi Help Center*, March 20, 2026. https://help.kalshi.com/en/articles/14026044-can-i-trade-on-kalshi-from-outside-the-united-states.

KuCoin. 2025. "Polymarket Profitability Crisis: 70% of Users Lose Money, Profits Concentrated among 0.04%." December 29, 2025. https://www.kucoin.com/news/flash/polymarket-profitability-crisis-70-of-users-lose-money-profits-concentrated-among-0-04.

Lambur, H., A. Lu, and R. Cai. 2020. "UMA Data Verification Mechanism: Adding Economic Guarantees to Blockchain Oracles." Risk Labs Technical Report, April 2020. https://github.com/UMAprotocol/whitepaper/blob/master/UMA-DVM-oracle-whitepaper.pdf.

Metzger, B. 2026. "Kalshi CEO Says Insider Trading Is 'Committing a Financial Crime' as He Backs New Bill." *Yahoo Finance*, January 7, 2026. https://finance.yahoo.com/news/kalshi-backs-bill-bar-prediction-165252701.html.

Ng, H., L. Peng, Y. Tao, and D. Zhou. 2026. "Price Discovery and Trading in Modern Prediction Markets." SSRN Working Paper No. 5331995. https://doi.org/10.2139/ssrn.5331995.

NPR. 2026. "Meta Plans to Release AI-Powered Prediction Market App, Documents Show." June 24, 2026. https://www.npr.org/2026/06/24/nx-s1-5869486/meta-prediction-market-app-ai.

Office of Senator Chris Murphy. 2026. "Murphy, Casar Introduce Bicameral Bill to Ban Prediction Markets on Government Actions, War, and Events Ripe for Rigging." March 17, 2026. https://www.murphy.senate.gov/newsroom/press-releases/murphy-casar-introduce-bicameral-bill-to-ban-prediction-markets-on-government-actions-war-and-events-ripe-for-rigging.

Pew Research Center. 2026. "Trading Volume on Prediction Markets Has Soared in Recent Months." May 27, 2026. https://www.pewresearch.org/short-reads/2026/05/27/trading-volume-on-prediction-markets-has-soared-in-recent-months/.

Polymarket Documentation. n.d. "Overview." Accessed September 5, 2026. https://docs.polymarket.com/trading/overview.

Polymarket Documentation. n.d. "Resolution." Accessed September 5, 2026. https://docs.polymarket.com/concepts/resolution.

Polymarket. 2026. "Polymarket Publishes Enhanced Market Integrity Rules Across Its DeFi Platform and CFTC-Regulated U.S. Exchange." *Business Wire*, March 23, 2026. https://www.businesswire.com/news/home/20260320997513/en/Polymarket-Publishes-Enhanced-Market-Integrity-Rules-Across-Its-DeFi-Platform-and-CFTC-Regulated-U.S.-Exchange.

Polymarket. 2026. "Polymarket Selects Chainalysis to Deploy First-of-Its-Kind On-Chain Market Integrity Solution." *Business Wire*, April 30, 2026. https://www.businesswire.com/news/home/20260430726176/en/Polymarket-Selects-Chainalysis-to-Deploy-First-of-Its-Kind-On-Chain-Market-Integrity-Solution.

Reichenbach, F., and M. Walther. 2025. "Exploring Decentralized Prediction Markets: Accuracy, Skill, and Bias on Polymarket." SSRN Working Paper No. 5910522. https://ssrn.com/abstract=5910522.

Robinhood. 2025. "Robinhood Launches Prediction Markets Hub." March 17, 2025. https://robinhood.com/us/en/newsroom/robinhood-prediction-markets-hub/.

Saguillo, O., V. Ghafouri, L. Kiffer, and G. Suarez-Tangil. 2025. "Unravelling the Probabilistic Forest: Arbitrage in Prediction Markets." In *7th Conference on Advances in Financial Technologies (AFT 2025)*, edited by Z. Avarikioti and N. Christin, vol. 354, 27:1–27:24. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. https://doi.org/10.4230/LIPIcs.AFT.2025.27.

Schapiro, M. L. 2010. "Testimony Concerning the Severe Market Disruption on May 6, 2010." U.S. Securities and Exchange Commission, May 11, 2010. https://www.sec.gov/news/testimony/2010/ts051110mls.htm.

Sobrado, B. 2025. "How Prediction Markets Actually Grew in 2025." *Forbes*, December 16, 2025. https://www.forbes.com/sites/boazsobrado/2025/12/16/how-prediction-markets-actually-grew-in-2025/.

Steiner, C. 2010. "Wall Street's Speed War." *Forbes*, September 9, 2010. https://www.forbes.com/forbes/2010/0927/outfront-netscape-jim-barksdale-daniel-spivey-wall-street-speed-war.html.

The Block. 2025. "Polymarket's Huge Year: $9 Billion in Volume and 314,000 Active Traders Redefine Prediction Markets." January 3, 2025. https://www.theblock.co/post/333050/polymarkets-huge-year-9-billion-in-volume-and-314000-active-traders-redefine-prediction-markets.

The Times of Israel. 2026. "Two Indicted for Using Classified Info to Place Online Bets on Military Operations." February 12, 2026. https://www.timesofisrael.com/two-indicted-for-using-classified-info-to-place-online-bets-on-military-operations/.

"Trader Made Nearly $1 Million on Polymarket with Remarkably Accurate Iran Bets." 2026. *CNN Politics*, March 24, 2026. https://www.cnn.com/2026/03/24/politics/iran-war-bets-prediction-markets.

Tsang, K. P., and Z. Yang. 2026. "The Anatomy of a Blockchain Prediction Market: Polymarket in the 2024 U.S. Presidential Election." arXiv:2603.03136. https://arxiv.org/abs/2603.03136.

Tsekova, D. 2026. "Spotify Challenges Prediction Markets After Song Chart Rigging." *Bloomberg*, July 2, 2026. https://www.bloomberg.com/news/articles/2026-07-02/spotify-challenges-prediction-markets-after-song-chart-rigging.

U.S. Department of Justice. 2026. "U.S. Soldier Charged with Using Classified Information to Profit from Prediction Market Bets." April 23, 2026. https://www.justice.gov/opa/pr/us-soldier-charged-using-classified-information-profit-prediction-market-bets.

Yurcaba, J. 2026. "U.S. Soldier Charged with Insider Trading over Maduro's Ouster." *NPR*, April 23, 2026. https://www.npr.org/2026/04/23/nx-s1-5797957/maduro-raid-charges-polymarket-insider.

Crypto Prediction Markets:

Collective Wisdom or a Hunting Ground for the Few?

Researcher

Seyeon Jo csy9512@naver.com

Corresponding Author

Taemin Oh

Editor

Sanghyun Park