Who Wins and Who Loses In Prediction Markets?*

Pat Akey[1], Vincent Grégoire[2], Nicolas Harvie[3], Charles Martineau[3]

[1]ESSEC Business School, CEPR, ECGI
[2]HEC Montréal
[3]University of Toronto

April 28, 2026

We study trading gains and losses on Polymarket, the world's largest prediction market platform. Our comprehensive dataset spans 2022 to 2026, covering more than 2.4 million users, $67 billion in volume, and 588 million trades. We document a striking profit concentration: the top 1% of users capture 76.5% of all trading gains. Despite this concentration, prices are well-calibrated and aggregate information efficiently. Gains flow almost entirely to sophisticated traders who trade using limit orders at prices that are advantageous relative to realized outcomes. Losses, by contrast, are associated primarily with liquidity-taking; longshot betting (trading at extreme prices) is a descriptive marker of losers but plays a smaller role once we control for activity scale. For roughly one in five losers, the lower-bound cost of taking liquidity alone is enough to flip their PnL from negative to positive. Monthly performance is only weakly persistent, however, so this cross-sectional pattern may reflect heterogeneous trading conditions or selection rather than persistent skill. Our results suggest that the informational benefits of prediction markets come at a cost to unsophisticated participants.

JEL Classification: G12, G14, G15, D47, D82.
Keywords: price efficiency, Polymarket, prediction markets, profit concentration, market microstructure.

*We thank Andreas Park and Anthony Sanford for helpful comments and suggestions and seminar participants at FEN Universidad de Chile, Pontificia Universidad Catolica de Chile, UTSC AI Management Talks. This research was undertaken, in part, thanks to funding from the Canada Research Chairs Program and from the Social Sciences and Humanities Research Council. Akey: ESSEC Business School, 3 avenue Bernard Hirsch, 95021 Cergy-Pontoise France (akey@essec.edu). Grégoire: HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine, Montréal QC, Canada H3T 2A7 (vincent.gregoire@hec.ca). Harvie: Rotman School of Management, University of Toronto, 105 St-George, Toronto ON, Canada, M5S 3E6 (nicolas.harvie@rotman.utoronto.ca) Martineau: Rotman School of Management and UTSC Management, University of Toronto, 105 St-George, Toronto ON, Canada, M5S 3E6 (charles.martineau@rotman.utoronto.ca).

1. Introduction

Prediction markets have exploded in popularity. Platforms including Polymarket and Kalshi allow users to trade on the outcomes of real-world events spanning politics, economics, sports, and entertainment. While prediction markets are not new, innovations in fintech and blockchain technologies have allowed these platforms to rapidly scale, with trading volume reaching $67 billion USD on Polymarket and $53 billion on Kalshi[1] in the last four years.

Media commentators and policy makers have raised questions about the gamification of trading and the role of prediction markets as an entry point into financial markets. Concerns have also been raised about "insider" trading on these platforms.[2] Prediction-market contracts are binary contingent claims that pay out on the resolution of an underlying event, and trading is zero-sum: every dollar gained by one trader is lost by another. Unlike equity markets, there is no equity premium, no dividends, and no passive investment options. Their short maturities naturally encourage short-horizon trading, which is precisely what the behavioral finance literature has suggested is most destructive to long-run wealth accumulation (Barber and Odean, 2000). Given the rapid growth of these platforms and their role as a potential entry point for young and retail investors, understanding who bears the losses has direct implications for investor welfare and market regulation. This raises a fundamental question: who wins and who loses in prediction markets?

We address this question using the complete transaction history of Polymarket from November 11, 2022 through March 29, 2026, covering over 2.4 million users and $67 billion in trading volume. Our main findings are fivefold. First, profits are concentrated: the top 0.1% of most profitable users capture 51.2% of all gains and the top 1% capture 76.5%. Meanwhile, 68.8% of users lose money. Second, Polymarket prices are well-calibrated: a contract priced at p resolves in its favor

[1] Statistic retrieved from `https://www.kalshidata.com/`

[2] See, for example, `https://www.cnbc.com/2025/09/23/young-people-crypto-meme-stocks-options-risk-financial-nihilism.html` and `https://www.reuters.com/legal/government/us-will-punish-fraud-insider-trading-derivatives-regulator-tells-congress-2026-04-16/`.

approximately p percent of the time, confirming that the market aggregates information efficiently. Third, the excess hit rate, defined as the average of realized outcomes minus trade prices, declines sharply through the PnL distribution, from roughly 15 percentage points in the top three earner percentiles to roughly -13 percentage points in the worst-performing percentiles. This finding suggests that profitability may be associated with some skill in identifying mispriced contracts. However, we urge caution in interpreting this finding as representing skill (or information) since we lack the tools typically used to assess performance in financial markets: a benchmark model of expected returns and a time series of the length typically used to assess performance, for example the performance of mutual fund managers. Fourth, descriptively, users who lose money trade considerably more often at extreme prices (below 10¢ or above 90¢) than users who gain: the bottom 95% of users place 56% of their trades at these prices, against 28% for the top 0.1% of earners. Once we control for activity scale and other behavioral characteristics, however, the conditional effect of extreme-price exposure is much smaller, indicating that longshot betting is a marker of losing rather than the primary mechanism of losses. Fifth, the strongest predictor of performance is liquidity provision: in the cross-section, a one-standard-deviation increase in a user's maker volume share is associated with a 9.3 percentage-point lower probability of loss, evaluated at the sample mean. This maker–taker asymmetry mirrors the findings of Barber, Lee, Liu, and Odean (2009) in Taiwanese equity markets and suggests that investor sophistication is an important predictor of the cross-sectional distribution of winners and losers on Polymarket.

These findings carry implications for both researchers and policymakers. For researchers, prediction markets offer a unique laboratory: the full population of trades, prices, and outcomes is publicly observable, making them a powerful setting for testing theories of market efficiency (see Thaler and Ziemba, 1988; Gomez Cram, Guo, Jensen, and Kung, 2025), skill persistence, and behavioral bias at scale. Do prediction markets function primarily as information-aggregation mechanisms or as speculative financial instruments? Our findings suggest both: price efficiency and extreme distributional inequality coexist on Polymarket. Prices closely track realized probabilities, yet a small group of sophisticated traders captures the vast majority of gains.

For policymakers, the answer has direct consequences for whether prediction markets warrant consumer-protection oversight. The extreme concentration of gains alongside the losses borne by most traders mirrors ongoing debates about retail options trading (Bryzgalova, Pavlova, and Sikorskaya, 2023) and sports betting (Baker, Balthrop, Johnson, Kotter, and Pisciotta, 2024). Prediction markets share key features with these activities: they are easily accessible, continuously available, and offer contracts that resolve quickly at binary payoffs. A growing body of evidence documents the household-level harms of such products. Access to gambling displaces household saving and investment (e.g., Cookson, 2018), and the legalization of sports betting causes a persistent deterioration in the financial stability of vulnerable households, raising overdraft rates, missed bill payments, and debt delinquency (Baker, Balthrop, Johnson, Kotter, and Pisciotta, 2024). Consistent with these financial harms, Ayers, Yeola, Allen, Desai, Poliak, Yang, and Smith (2025) document a 23% rise in gambling-addiction help-seeking following sports-betting legalization. Beyond these aggregate harms, Packin and Rabinovitz (2026) argue that prediction-market platforms constitute a public-health threat, employing design elements and nudges similar to online gambling platforms that exploit cognitive biases and encourage further gambling (Newall, 2025).

Prediction markets have long been studied as mechanisms for aggregating dispersed information into prices. Hayek (1945) first argued that prices coordinate decentralized knowledge more efficiently than any central planner, and prediction markets represent a direct institutional embodiment of this idea. Forsythe, Nelson, Neumann, and Wright (1992) provided early empirical validation using the Iowa Electronic Markets (IEM), showing that real-money prediction markets outperformed polls in forecasting US election outcomes. Wolfers and Zitzewitz (2004) offer the canonical survey of the field, documenting the conditions under which prediction market prices can be interpreted as physical probabilities and reviewing evidence on their forecasting accuracy across political, economic, and sporting events. Manski (2006) raises an important cautionary note, showing that the mapping from prices to probabilities depends on assumptions about trader risk preferences that are rarely tested, a limitation that applies to our setting as well.

A central finding in the prediction market literature is that market accuracy does not require individual rationality. Forsythe, Nelson, Neumann, and Wright (1992) introduced the marginal trader hypothesis: a small group of active, well-informed traders who set limit orders drive prices to efficiency even when average participants exhibit partisan bias and wishful thinking. Oliven and Rietz (2004) examine IEM trader-level data and show that while 37.7% of price-taker orders violated no-arbitrage conditions, market makers violated them only 5.4% of the time. The closest intellectual antecedent in financial markets is Barber, Lee, Liu, and Odean (2009), who use brokerage data to show that a small number of sophisticated traders systematically profit at the expense of retail participants, a finding we complement with our results in the prediction-market context. This division between price-setting makers and price-taking participants foreshadows the maker-taker asymmetry we document on Polymarket, where liquidity provision is the strongest cross-sectional predictor of profits.

A second strand of literature documents behavioral anomalies and efficiency failures in prediction markets. Snowberg, Wolfers, and Zitzewitz (2012) document the favorite-longshot bias: the tendency for low-probability outcomes to be overpriced and high-probability outcomes to be underpriced. Snowberg and Wolfers (2010) provide evidence that this bias is driven by probability misperceptions consistent with prospect theory's weighting function. This bias is directly relevant to our setting, as users on Polymarket disproportionately trade contracts at extreme prices. More broadly, Barber and Odean (2000) show that individual stock investors lose money through excessive trading, a pattern that resonates with our findings on losses concentrated among heavy traders of longshot contracts.

Polymarket's on-chain transparency has recently enabled a new wave of prediction market research. Tsang and Yang (2025) provide a comprehensive transaction-level analysis of Polymarket's 2024 U.S. election market, documenting key liquidity episodes and showing that Kyle's λ declined as market liquidity matured. Reichenbach and Walther (2025) identify persistent skill differentials among Polymarket traders, with market prices closely tracking realized probabilities. Becker (2025), analyzing over 72 million trades on Kalshi, documents persistent wealth transfer from takers to makers and a pronounced asymmetry in which takers disproportionately favored affirmative bets at longshot prices. Sirolly, Ma, Kanoria, and Sethi (2025) develop network-based wash trading detection methods and estimate that approximately 25% of historical Polymarket volume may be artificial, peaking at 60% during December 2024. Gomez Cram, Guo, Jensen, and Kung (2025) construct a high-frequency measure of earnings expectations from Polymarket prices, showing that market-implied expectations are more accurate, less biased, and lead in price discovery relative to analyst forecasts, providing direct evidence that prediction markets aggregate fundamental information efficiently in financial contexts. Another strand of research finds that prediction market prices provide information comparable to that of professional forecasters (Diercks, Katz, and Wright, 2026), yet the traders and the trading behavior that generate these prices remain poorly understood.

Our contribution is distinct from and complementary to these contemporaneous studies. Becker (2025) documents maker-taker wealth transfer in aggregate on a centralized exchange (Kalshi), whereas we provide the first wallet-level accounting of user trading profits and losses on a decentralized platform with full transaction transparency. Reichenbach and Walther (2025) study skill persistence and forecasting accuracy. In contrast, we document the distributional consequences: the concentration of gains among a small group of traders, the losses borne by the majority of traders, and the role of maker-taker dynamics in driving this redistribution. To our knowledge, no study documents the full distribution of profits and losses across participant types in any prediction market.

Our accounting is complete at the level of the reconciled on-chain trade stream, but is subject to several known limitations. First, users operating multiple wallets appear as separate accounts. As we describe in more detail in Section 3.2, this is unlikely to impact the main conclusions of our paper unless there are many such users and such users are systematically trading in a completely different manner in their different wallets. Second, there are concerns that some users may engage in "wash trading," whereby different accounts engage in correlated trading in an effort to drive trading volume or attention in different markets. Our robustness checks suggest that wash trading is unlikely to be driving our main findings. Finally, we do not observe off-platform trading, which may be relevant for some users, especially sophisticated traders who may be engaging in arbitrage across platforms. However, we show that the majority of trading volume on Polymarket is not easily exploitable by arbitrageurs.

2. Institutional Background

Polymarket and Kalshi are decentralized prediction market platforms where users trade binary contracts on real-world outcomes. Our analysis focuses on Polymarket, which was launched in 2020 and has grown into the largest prediction market by volume. Unlike traditional financial exchanges, it operates without a central intermediary: trades are executed and settled automatically via smart contracts on the Polygon blockchain. In January 2022, the CFTC charged Polymarket's operator (Blockratize, Inc.) with offering unregistered binary options to the public and imposed a $1.4 million civil monetary penalty, after which the platform restricted US nationals from participating.

Every Polymarket event is structured as a set of at least one market (binary contracts), one for each possible outcome.[3] For example, an event such as "2026 NHL Stanley Cup Champion" would have separate markets for each team. One such market would be "Montreal Canadiens", with two possible outcomes, Yes and No, that each trade as a distinct token. Contracts are denominated in US dollar coin (USDC), a dollar-pegged stablecoin.

Holding a Yes token entitles the owner to \$1 if the outcome occurs and \$0 otherwise; holding a No token entitles the owner to the complement.[4] Settlement is automatic: when a contract resolves (see below for more details on contract resolution), the blockchain smart contract attributes the value of \$1 to winning positions and \$0 to losing ones. Users are then required to claim the winning positions for USDC. This removes counterparty risk and ensures complete settlement transparency.

Users can trade these tokens at any time before expiration or exchange a pair of complementary tokens for 1 USDC. Polymarket uses a hybrid architecture combining a central limit order book (CLOB) matching engine with on-chain settlement of conditional tokens.[5] Users can either post limit orders (makers) or trade against existing liquidity (takers). Trades are executed on the Polygon blockchain, a low-cost Ethereum-compatible network. Each user is identified by a unique wallet address rather than a name or account number. Every transaction is permanently and publicly recorded on-chain, giving participants and researchers an unedited view of all trading activity since the platform's switch to Polygon in late 2022. We use this record of trades to calculate user-level profits and losses.

Prices in prediction markets are often referred to as "probabilities". Under the assumption of risk neutrality and ignoring time-value of money, the contract price at any moment reflects the

[3]The actual contract structure allows for markets with more than two outcomes tokens. In practice, only 39 of the 729,133 markets in our dataset have more than two outcomes. We include these markets in our profitability analysis but exclude them from our analysis of calibration and accuracy.

[4]Not all binary markets are structured as Yes/No pairs. For example, a market on the winner of a sporting event with two opposing teams might have a pair of tokens, one for each team. These are just labels, and the underlying economics are the same as a Yes/No pair.

[5]The Polymarket exchange smart contracts seamlessly aggregate liquidity across all outcomes of a given market, allowing users to easily trade without needing to worry about the underlying order book structure. For example, the ask side of the book for the Yes token is a mirror image of the offer side for the No token, and vice versa.

market's implied probability of the event occurring. A contract trading at \$0.35 means the market assigns a 35% probability to the outcome. Contracts that users can bet on span a wide range of categories: politics, sports, economics, and geopolitical events, with expiration dates ranging from days to months or upon realization of the event.

Polymarket did not charge trading fees on most markets during our sample period. Therefore, users who post limit orders (makers) earn the bid-ask spread as compensation for providing liquidity, while users who trade against existing liquidity (takers) pay the spread as a transaction cost. The minimum tick size is \$0.01, except for markets with a price below \$0.04 or above \$0.96, which have a tick size of \$0.001.[6] However, certain markets have taker fees to fund the Maker Rebates Program and to incentivize deeper liquidity and tighter spreads.[7] We remove these markets from our analysis in robustness checks and find that our main results are unchanged.

Contract resolution on Polymarket is determined by an oracle, a decentralized mechanism that reports real-world outcomes to the blockchain.[8] Once a contract expires, the outcome is proposed by a designated proposer and enters a challenge window during which any participant can dispute the resolution. To initiate a dispute, the challenger must post a bond in UMA tokens (currently equivalent to several hundred to a few thousand US dollars depending on the contract) to deter against frivolous challenges, ensuring that disputes are only initiated when the challenger has genuine conviction that the proposed resolution is incorrect. The dispute is then adjudicated by UMA token holders who vote on the correct outcome, with the losing side forfeiting their bond to the winning side.[9]

[6]Note that this is not defined as a unique pricing grid with fixed tick sizes as common in equity markets. Instead, once the price of a contract crosses the \$0.04 or \$0.96 threshold, the market switches to a finer grid until resolution, so we can observe prices at sub-\$0.01 increments inside the \$0.04–\$0.96 range. This also means that it is not possible to post limit order at sub-\$0.01 increments, even for extreme prices, until the market has switched to the finer grid.

[7]The following market types have fees enabled: All crypto markets (1 hour, 4 hours, daily, and weekly starting March 6, 2026, for new markets), National Collegiate Athletic Association (NCAA) basketball games (starting February 18, 2026, for new markets), Serie A (football) markets (starting February 18, 2026, for new markets).

[8]Polymarket uses the Universal Market Access (UMA) Protocol as its resolution oracle.

[9]For example, contested resolutions may arise in contracts with ambiguously worded conditions, such as those tied to qualitative political or geopolitical events where the resolution criteria leave room for interpretation.

3. Data

Polymarket operates on the Polygon blockchain, where order settlements are recorded on-chain. We collect the complete transaction history by combining two complementary data sources. The events and markets metadata is from Polymarket's Gamma and CLOB APIs. Trade information is obtained from the Polymarket SubGraph by Goldsky, a GraphQL API that allows for direct access to Polygon's on-chain order fill events. These events provide a record of all transactions, including maker and taker wallet addresses for each settlement. We link these two sources to construct a unified dataset that associates every trade with its price, quantity, direction, counterparty identities, and contract outcome, yielding the complete population of trades on the platform.[10] Our sample spans November 11, 2022 through March 29, 2026; we end on that date because Polymarket introduced trading fees on all markets on March 30, 2026, which materially changes the economics of trading on the platform. For each trade, we observe the wallet addresses of the maker and taker, the contract identifier, timestamp, direction (buy/sell), quantity, price, and the contract resolution outcome for markets that have resolved. The full list of data sources, the procedure that reconciles raw on-chain events into clean trade records, and the construction of our main dataset are described in Appendix A.

3.1. Trade data construction

The raw trade data is first standardized into a single dataset which contains the recorded information for each trade across all Polymarket markets, including the traded token, price (in USDC per share), quantity (in shares), maker and taker addresses, and outcome metadata. Almost all the markets in our sample are binary markets, i.e., those with exactly two mutually exclusive outcomes (e.g.,

[10] Polymarket's Gamma API also allows downloading trades, but this approach yields an incomplete sample of trades because it limits the number of trades per market that can be retrieved. Direct querying of Polygon's Blockchain is necessary to obtain a complete sample.

Yes/No, or Team A vs. Team B). For those, we normalize all trades to express prices from the perspective of a single reference outcome, so that a price of p always represents the market-implied probability of that outcome. Specifically, we restrict to Yes/No markets and normalize to the "Yes" token: trades on the "No" token are converted by replacing the token identifier with its "Yes" counterpart and inverting the price to $1 - p$. We normalize all two-outcome markets by normalizing to the first outcome by index, enabling accuracy analysis on non-Yes/No binary markets such as head-to-head matchups. We describe the procedure that we use to identify the winning outcome for each market in Appendix A.3.

3.2. Constructing User Accounts and PnL

Our analysis focuses on wallet-level trading activity and profitability. We construct user accounts by aggregating all trades associated with the same wallet address. For simplicity, we refer to each wallet address as a "user", though we acknowledge that some traders may operate multiple wallets, which would cause us to undercount their activity and overstate the number of distinct users.

At each daily snapshot at midnight UTC time, each user's token holdings are marked to market by joining their positions to the most recent available price or the resolution price, if it is known at that time. The portfolio value is then the sum of position $\times$ price across all token holdings, and PnL $=$ portfolio value $+$ net USDC balance, meaning the USDC balance tracks realized gains and losses while the portfolio value tracks unrealized ones.[11] Note that we do not observe users' USDC balances. Instead, we infer their net balances from their transactions. For example, a user that buys 100 Yes tokens for \$ 0.63 each will have a net USDC balance of -\$63 and a positive position in the token. Positions are forward-filled from a user's first trade through March 29, 2026, so every

[11] Our PnL definition does not net out Polygon gas costs. Polymarket sponsors gas fees for users trading through the official website, so these costs are zero for the typical participant. Sophisticated traders who interact directly with the smart contracts on the Polygon blockchain do incur per-transaction gas fees, but these are small in absolute terms.

user has a daily snapshot for their entire active life. The full construction of the position and PnL panel, including the treatment of market resolution and the marking-to-market logic, is described in Appendix A.4.

To validate our PnL calculations, we selected a set of users with a minimum trade count of 10, traders whose last trade occurred strictly before a cutoff date (March 29, 2026), for whom our local trade data is guaranteed to be complete. From these eligible users, we randomly selected 1000 using stratified sampling by trading activity (low: <100 trades, medium: 100–1,000, high: >1,000) and sampled proportionally (40%/40%/20%). For the resulting list of wallet addresses, we then cross-checked their PnL against the live Polymarket Data API, and found a perfect match between our calculations and the API data.

3.3. Descriptive Statistics

Figure 1 illustrates Polymarket's growth over our sample period. From the platform's launch through early 2024, user adoption was modest, with roughly 1,000 new users joining per month and monthly trading volume hovering around $1 million. The platform's growth then accelerated sharply in the lead-up to the 2024 U.S. presidential election: by November 2024, monthly new user registrations reached approximately 100,000 and monthly trading volume surged to around $1 billion. At the end of our sample, the monthly trading volume exceeded $10 billion.

Table 1 summarizes key trading characteristics across all users in our sample. Trading activity is highly skewed: the median user executes 28 trades across 10 distinct markets (questions), while the mean is substantially higher (474 trades, 54 markets), reflecting a long right tail of highly active participants. Most users trade exclusively as takers—the median fraction of maker volume is 0%, and the 75th percentile reaches only 39.5%—though the mean of 17.9% indicates a nontrivial minority that supplies liquidity. PnL is also highly dispersed: the median user loses $2, consistent with the aggregate zero-sum nature of prediction markets net of fees, but individual outcomes range from losses of over $10 million to gains exceeding $22 million. Yet, the 95th percentile for PnL of $252 suggests that the vast majority of users experience modest outcomes, with extreme profits and losses concentrated among a small fraction of participants.

Figure 2 shows the monthly count of new events and markets launched on Polymarket. The platform's early period was characterized by a steady stream of new markets, averaging around a little more than 100 markets per month. However, from mid-2024, there was a dramatic surge in market creation, peaking at over 100,000 new markets in 2026.

Table 2 breaks down the distribution of events and markets and additional statistics by category. We assign each market to one of seven categories: Sports, Crypto, Finance, Politics, Tech, Culture, and Weather.[12] Table 2 reports summary statistics by category. Sports dominates by number of markets (321,231; 52% of all markets) and user reach (1.57 million users have participated at least once in this category, or 63% of users), and accounts for 37.9% of total volume. Crypto has the highest trade count (394 million trades, 67% of all trades) driven by a large number of markets (206,491), but its per-market average volume ($61,804) is far lower than Politics ($705,553), which generates 31.5% of total volume from only 29,926 markets. Culture, Finance, Tech, and Weather together account for less than 12% of total volume.

4. Who Wins and Who Loses?

In this section, we analyze the distribution of trading outcomes on Polymarket to identify who the winners and losers are, and what characteristics predict profitability.

[12] This classification relies primarily on Polymarket's platform-assigned event tags, with fallback rules for events and orphan markets that the tag-based procedure does not classify. Appendix A.2 details the full assignment procedure.

4.1. Concentration of Gains and Losses

Figure 3 shows the distribution of trading gains and losses across Polymarket users. The x-axis ranks users with positive PnL in green and with negative PnL in red, by their percentile within that group, while the y-axis shows the cumulative share of total gains and losses they account for. Across our sample of 766,188 users with positive PnL, the aggregate net gain of winners totals $1.1 billion.

These gains are strikingly concentrated among a small number of individuals. The top 0.1% of users with positive PnL alone captured 51.2% of all gains, while the top 1% captured 76.5%, the top 5% captured 91.4%, and the top 25% accounted for 99.2%. The vast majority of winning users, therefore, earned only a negligible share of total profits. This extreme concentration is consistent with a tiny fraction of traders accounting for nearly all of the net surplus, with the much larger pool of losing participants on the other side. Focusing on dollar profits rather than percentage returns in this analysis is deliberate: Cziraki and Gider (2021) show that the correlation between dollar profits and percentage returns is moderate, because returns are negatively correlated with trade size and frequency, so dollar profits better identify genuinely skilled and informed investors who express their edge through trade size, not just direction. Figure 3 also shows that losses are also concentrated but to a lesser extent: the top 0.1% of losers account for 43.7% of total losses, the top 1% for 68.1%, the top 5% for 86.1%, and the top 25% for 98.3%.

Figure 4 tracks this concentration over time. Prior to the 2024 U.S. presidential election, the top 1% of winners captured approximately 40–50% of total gains. In the lead-up to the election, an event that spurred the growth of users participating in prediction markets, the concentration of gains among the top 1% rose to approximately 80%. A similar pattern holds for losses. This elevated concentration has remained relatively stable through the end of our sample period even as the total volume grew fourfold.

4.2. How efficient is Polymarket?

A fundamental question is whether Polymarket prices accurately reflect the true probabilities of outcomes, i.e., whether they are well-calibrated. Panel A of Figure 5 addresses this directly by plotting the realized frequency of winning outcomes against price bins. If markets are efficient and there is no risk premium, contracts trading at price p should resolve in favor of the relevant outcome approximately p percent of the time, tracing out the 45-degree line. Across the full price range, Polymarket prices closely track this benchmark. Contracts in the 10–20¢ range resolve in favor of the outcome roughly 10–20% of the time, contracts near 50¢ resolve at close to a coin flip, and high-priced contracts near 80–90¢ resolve favorably at correspondingly high rates. This aggregate calibration is driven primarily by the high-volume tail of markets; calibration in low-volume markets is noisier, and we provide the full breakdown by volume quintile and by fixed volume bucket in Figures C3 and C4 of the appendix.

The remaining panels of Figure 5 examine price efficiency across market categories, replicating the binned calibration plot for Politics, Sports, Crypto, and others. This allows us to assess whether the efficiency shown in Panel A is consistent across all categories or driven by a subset of particularly liquid or well-covered markets. The figure shows that for Sports, the 7-day and 1-day horizons track the 45-degree line closely. For Crypto, Politics, and Finance, prices become more efficient as they approach expiration, with the 1-day horizon converging closer to the 45-degree line. The Tech, Culture, and Weather categories show more pronounced deviations from the 45-degree line, even at the 1-day horizon, suggesting that these markets may contain more mispricings that skilled traders could potentially exploit.

4.3. Do profitable traders have better forecasting skill?

We next examine the relationship between user profitability and forecasting skill. We first define a measure of skill, the excess hit rate (EHR), which we compute for each user i as:

$$\mathrm{EHR}_i = \frac{1}{K_i}\sum_k (o_k - p_k), \tag{1}$$

where K_i is the number of trades k for user i, every trade is recast to the buy-side perspective so that $o_k \in \{0, 1\}$ equals 1 if and only if the contract *purchased* in trade k resolves in-the-money, and p_k is the buy-side normalized trade price.[13] We also compute a volume-weighted version of this metric, $\mathrm{EHR}_{\mathrm{vw},i} = \sum_k q_k(o_k - p_k)/\sum_k q_k$, where q_k is the traded quantity in shares.[14] Weighting by volume ensures that large trades impact the EHR more. A positive EHR indicates that a user's trades resolve more favorably than their prices suggested; a negative EHR indicates the opposite.

Figure 6 plots EHR and $\mathrm{EHR}_{\mathrm{vw}}$ across the full PnL distribution. We rank users by PnL and assign them to 200 half-percentile bins from the most profitable (percentile 0) to the least profitable (percentile 100). For each bin we plot the mean (solid line, with 95% confidence band) and the median (dashed line) of each metric. Vertical dotted lines mark the bins containing the PnL thresholds of \$100, \$0, and −\$100, which fall near percentiles 8, 30, and 85, respectively, and together with the figure make clear that roughly the top 8% of users earn more than \$100 in profit, a further 22% earn between \$0 and \$100, and the rest lose money, with about 15% of users losing more than \$100.

[13] Buy-side normalization recasts every trade as a purchase: a sell of outcome X at price p is treated as a buy of the complementary contract at $1 - p$. This gives every user position a single sign convention and lets us interpret $o_k - p_k$ uniformly across both sides of the market. Note that Polymarket maintains a consolidated order book for each market, so a buy limit order for a Yes contract at price p also appears as a sell limit order for the No contract at price $1 - p$ in the order book, and vice versa; Polymarket seamlessly matches buy and sell orders across the two sides of the market.

[14] Under the matched-notional convention (each share traded represents \$1 of volume in USDC, since a Yes-buyer and a No-buyer together pay \$1 per share), share-weighting and USDC-weighting coincide; we use "volume-weighted" and "share-weighted" interchangeably for $\mathrm{EHR}_{\mathrm{vw}}$.

The top panel shows that the mean EHR is strongly positive at the top of the PnL distribution, starting near $+0.15$ at percentile 1 and declining steadily to around $+0.02$ by the time cumulative profit falls below \$100 (percentile 8). Mean EHR continues to decline through the winner cohort, reaches zero around the \$0 PnL threshold, and then turns negative throughout the negative-PnL side of the distribution, fluctuating between roughly -0.05 and -0.15. The loss tail is noisier than the gain tail because the number of resolved trades per user shrinks, but the mean remains consistently below zero across every loser bin. The median EHR (dashed line) tracks the mean's sign but is always smaller in magnitude—about $+0.05$ to $+0.10$ at the top, near zero in the middle, and -0.03 to -0.07 in the loser tail—reflecting the long-tailed distribution of per-user EHR: a minority of traders in each bin drive the extremes of the mean. The bottom panel reports the same picture for the volume-weighted metric, where magnitudes are uniformly larger than in the top panel: mean EHR_{vw} reaches about $+0.20$ for the top percentiles and descends to roughly -0.15 to -0.18 in the loser tail. The larger magnitudes indicate that both winners and losers tilt their largest trades further in the favorable (respectively unfavorable) direction than their average trade.

Overall, these results indicate a strong positive relationship between profitability and forecasting skill: the most profitable traders consistently buy contracts that resolve more favorably than their prices implied, while the least profitable traders consistently buy contracts that resolve less favorably. However, we urge caution in interpreting this finding as representing skill (or information) since we lack the tools typically used to assess performance in financial markets: a benchmark model of expected returns and a time series of the length typically used to assess performance, for example the performance of mutual fund managers.

4.4. Is performance persistent?

We next examine whether monthly performance is persistent. Having established that cumulative PnL is concentrated and that a small proportion of traders have positive EHR while most have negative EHR, we ask a different question: is individual performance consistently positive or negative across time? To answer this question, we classify traders on a monthly basis into five groups based on the PnL they realized during a given month,[15] with bins of <-$100, -$100 to -$10, -$10 to $10, $10 to $100, and >$100. We then examine in which group each user falls into 1, 3, and 6 months later.[16] We report the fraction of users that remain in their starting group, transition to another group, or drop out (no trades over the horizon month).

Panel A of Table 4 presents results of this analysis at the one-month horizon. We can see that lower performance positively predicts exit from Polymarket, although even the best performing traders exit. For example the one-month exit rate ranges from 44.1% for the worst performing traders to 28.6% for the highest performing traders. While there is some evidence for persistent *underperformance*, there is little evidence of persistent *overperformance* at the one-month frequency. For example, conditional on continuing to trade, the lowest performing group has a 37% chance of remaining in the lowest performing group (=0.208/(1-0.441)) and a 21.4% chance of moving to the highest performing group (=0.12/(1-0.441)), while the highest performing group has roughly equal odds of staying in the highest performing group (33.6% = 0.24/(1-0.286)) or moving to the lowest performing group (33.2% = 0.237/(1-0.286)). Interestingly, the "breakeven" traders (those with profits between -10 and 10) exhibit a strong likelihood of remaining in the same group. Indeed, conditional on trading, these traders exhibit an 89.2% (= 0.569/(1-0.362)) likelihood of having the

[15] Monthly mark-to-market PnL is the change in their PnL from the prior month's end. If a user did not trade in a given month, there is no user-month observation. Users that trade over many months will appear multiple times.

[16] Note these calculations examine the performance in individual months relative to the reference month, not cumulative performance over the time horizon. For example, for a user-month observation in January, the 6-month performance looks at the PnL in July, not the cumulative PnL from February to July.

same level of profits one period later. We observe qualitatively similar patterns across other horizons. Thus, while we find some evidence that a small group of traders have higher than "expected" hit rates, such performance does not seem to persist through time, which challenges the interpretation that many traders have substantial skill.

4.5. What Explains the Losses?

To identify determinants of negative trading outcomes, Table 5 estimates probit regressions where the dependent variable is an indicator equal to one if a user's mark-to-market PnL is negative. Reported values are marginal effects at the sample mean, so each cell gives the change in the probability of loss produced by a one-unit change in the regressor, evaluated at the mean of the covariates; standard errors are robust. We select a set of variables motivated by the market microstructure and behavioral finance literatures. *Frac Extreme Price* measures the share of a user's trades executed at prices below 10¢ or above 90¢, targeting contracts near the boundary of certainty whose payoffs are highly nonlinear and whose mispricings are difficult to exploit without an informational edge. *Frac Maker Volume* is the fraction of a user's total volume supplied as a passive limit order, distinguishing liquidity providers from takers; *Log N Trades* and *Log Total Volume* are natural logarithms of the number of executed trades and total USDC volume, respectively, controlling for activity scale; *Category HHI* is the Herfindahl-Hirschman Index of a user's volume concentration across market categories, where higher values indicate specialization in fewer categories; and *Counterparty HHI* is the analogous index across counterparties, which indicates whether a user trades with a wide variety of counterparties or concentrates their activity with a few.[17] Column (4) adds category fixed

[17] The intended interpretation of this counterparty concentration variable is that users who trade with few counterparties are trading with "frequent" traders, most likely market makers or other sophisticated traders. Another possibility is that users with high counterparty concentration are engaging in wash trading. Wash trading is the practice of executing trades primarily to move the price of an asset, rather than to gain a profit, by trading with oneself or a related entity. Throughout the main paper we include all users in every test, which would include potential wash traders. To assess their potential impact, we implement a simple wash-trading detection procedure. Appendix B describes our HHI-based identification procedure, reports summary statistics on flagged wallets and their volume, and replicates Section 4.1's concentration results and the probit specifications reported here with flagged wallets

effects: a binary indicator for each of the seven market categories equal to one if the user has any trading volume in that category, with Sports omitted as the reference. We deliberately exclude any performance-based measures, such as prediction accuracy or hit ratios, to isolate behavioral rather than outcome predictors of loss, making sure the regression captures ex-ante trading patterns rather than ex-post realizations. Table 5 reports the results for all users (columns 1–4), for users with more than 100 trades (columns 5–6), and for users with more than 1,000 trades (columns 7–8).

The dominant economic effect belongs to *Frac Maker Volume*. In the full specification for all users (column 4), a one-unit increase in the volume-weighted maker share is associated with a 34.4 percentage-point reduction in the probability of loss. Because few users are near either extreme of this share, we express its economic impact in standard-deviation units: a one-standard-deviation increase in *Frac Maker Volume* reduces the probability of loss by 9.3 percentage points in column (4), 9.5 pp in column (3), and 9.0 pp in column (6). This magnitude is the largest in the cross-section and dwarfs every other behavioral predictor; it shrinks to 4.4 pp per 1 SD in the $>$ 1,000-trade subsample (column 8), where the activity filter already absorbs much of the cross-sectional variation in maker share. The microstructure literature documents that informed investors often trade via limit orders rather than market orders (Kaniel and Liu, 2006; Brogaard, Hendershott, and Riordan, 2019), so the maker-share effect we estimate likely bundles a liquidity premium with superior information. As with the Section 4.3 caveat, we caution that this cross-sectional association does not establish a causal role for market-making per se.

The marginal effect of *Frac Extreme Price* is small in magnitude and varies across specifications. In the simplest specification for all users (column 1), the marginal effect is -1.2 pp. Adding activity-scale controls and category fixed effects shrinks it further, leaving only -0.3 pp in the fully specified column (4). Conditioning on more active users turns the marginal effect modestly

excluded. None of the paper's main findings are materially affected.

positive in columns (5)–(6) (+0.7 pp and +1.2 pp), and then sharply negative in the > 1,000-trade subsample (columns (7)–(8), −12.9 pp and −12.8 pp), consistent with a selection story: among users who persistently transact thousands of times at extreme prices, exposure to the long-shot segment no longer predicts losses in the same direction as in the broader population of casual traders. The substantive takeaway is that, once we control for how much and how often a user trades, extreme-price exposure has only a second-order direct association with the probability of loss in the all-user specification, even though descriptively losers trade at extreme prices roughly twice as often as the top earners (Table 6).

Log N Trades has a small negative marginal effect in column (2) (−0.1 pp) but turns positive once category controls enter, reaching +0.4 pp in column (3) and +1.7 pp in column (4). Among users with more than 100 trades the effect grows to +2.7 pp and +2.9 pp in columns (5)–(6), consistent with overtrading as a behavioral marker of overconfidence (Barber and Odean, 2000). *Log Total Volume* has a small positive effect in the full sample (column (4): +2.0 pp) but turns negative among users with more than 100 or 1,000 trades (−2.1 to −4.3 pp per log unit in columns 5–8), suggesting that conditional on trade count, users who commit larger amounts per trade are somewhat less likely to lose.[18] The *Log N Trades* marginal effect flips sign in the > 1,000-trade subsample (−0.4 and −0.6 pp in columns 7–8 versus −0.1 to +2.9 pp across columns (2)–(6)), consistent with a similar selection story: once we condition on users who have already executed more than a thousand trades, the within-subsample variation in activity no longer predicts losses in the same direction as in the broader population.

The *Category HHI* coefficient is informative about what specialization captures in this sample. In column (3), without category fixed effects, greater concentration in fewer categories is associated with a +8.1 pp higher probability of loss, a pattern consistent with casual traders who restrict their

[18] The two variables are moderately correlated (0.65) and together control for the scale of trading activity.

volume to a single topic without any genuine informational advantage. Once category fixed effects enter (column 4), however, the marginal effect flips to -5.5 pp: conditional on which categories a user actually trades in, higher within-category concentration instead predicts a lower probability of loss, consistent with specialization by choice among traders focused on markets they know well. The same direction emerges in the more active subsamples, where higher concentration predicts lower loss probabilities throughout (columns 5–8) and the addition of category fixed effects in column (6) deepens the effect from -1.0 to -6.7 pp.

Counterparty HHI displays an informative pattern that strengthens with activity. In the full sample the marginal effect changes sign across specifications (-0.5 pp in column (3) and $+3.6$ pp in column (4) once category fixed effects enter) and grows sharply among more active traders, reaching $+20.6$ pp for users with more than 100 trades (column 6) and $+49.4$ pp for those with more than 1,000 trades (column 8). A high counterparty HHI means a user repeatedly transacts with the same small set of counterparties—a pattern consistent with liquidity takers who repeatedly trade with the same market makers.

The category fixed effects in column (4) are themselves informative. Relative to users who trade only in Sports, users with any exposure to Crypto, Politics, Tech, Weather, or Finance have a 3.3 to 7.4 percentage-point lower probability of loss in the full sample (each coefficient is highly significant), while having any exposure to Culture is effectively neutral ($+0.2$ pp). Most of these coefficients shrink toward zero in the > 100-trade subsample (column 6), and in the $> 1{,}000$-trade subsample (column 8) exposure to Crypto ($+5.7$ pp) or Politics ($+4.1$ pp) flips to predicting a higher probability of loss. This reversal is consistent with the most committed segment of the user base differing systematically across categories: among users with thousands of trades, the Sports-only cohort appears to contain a relatively higher share of profitable traders than comparably active users who also bet on Crypto or Politics.

Taken together, the results show that losses are driven primarily by acting as a liquidity taker, and, to a lesser extent, by trading a large number of small trades and by concentrating volume on a small set of counterparties. Maker volume share is by far the largest behavioral predictor across every specification and subsample, consistent with the view that liquidity providers tend to be more profitable than liquidity takers. Once category fixed effects are included, within-category specialization is associated with lower loss probabilities, and users whose volume is confined to Sports are not distinctive performers in the full sample but fare relatively better among the most active traders.

Table 6 provides a descriptive complement to the probit results by directly comparing the trading characteristics of the most profitable users against the bottom 95%. The contrasts are striking and corroborate the probit marginal effects. The top 0.1% of earners supply 47.2% of their volume as maker, more than two and a half times the 17.1% maker-volume share of the bottom 95%, a gap of 30.2 percentage points that is highly significant. Their exposure to extreme-price contracts is almost half that of losing traders: 28.2% versus 56.4%, a difference of −28.2 percentage points. Top traders also exhibit substantially higher excess hit rates—0.097 for the top 0.1% against −0.027 for the bottom 95%—confirming that profitability is inseparable from the ability to identify mispriced contracts. The Category HHI gap is smaller in magnitude (3.7 percentage points for top 0.1% vs. bottom 95%) but statistically significant and positive, echoing the negative marginal effect of *Category HHI* in the full-specification probit: top traders are *more* concentrated by category, not less, once we condition on which categories they trade. Together, the two tables paint a consistent picture: winning traders are sophisticated liquidity providers, while losing traders are liquidity takers who repeatedly overpay for low-probability outcomes.

4.6. Spread decomposition and liquidity provision

We now examine the role of liquidity provision and spreads on trader profitability. We first note that we can observe, for each trade, which user provided or took liquidity. We begin by estimating a net-of-transaction-fee PnL by adjusting each observed transaction to improve the price for the liquidity taker by the minimum possible effective spread. For example, if a user buys a contract at 40¢ via market order, we compute a counterfactual PnL as if they had bought at 39.5¢ instead, which is the minimum possible effective half-spread in the standard tick regime. The counterfactual PnL for the liquidity provider will also use this adjusted price so that aggregate PnL still sums up to zero. [19] This adjustment gives us a lower bound on how much of the observed losses can be attributed to transaction costs rather than to poor forecasting. The true effect is likely much larger, especially in less liquid markets where the effective spread can be several times the minimum tick size.

Table 7 shows how much of the observed losses can be attributed to transaction costs rather than to poor forecasting. Because we remove only the tick-size floor rather than the actual effective spread, the resulting *Frac. Spread* is a lower bound on the share of realized gains or losses attributable to the spread. Among losers, the spread is a substantial driver of losses only for the smallest losers: for the bottom 0–20 percentile, even this lower-bound adjustment accounts for 195.8% of realized losses, meaning that these users would have been slightly profitable in the absence of spread costs. The fraction falls sharply as losses grow larger, to 59.8% at the 40–60 percentile and just 0.9% at the top 0.1% of losers, showing that the most severe losses are overwhelmingly driven by bad forecasting rather than transaction costs. At the user level, removing this lower-bound spread cost flips 18.5%

[19] The effective spread is the difference between the price at which a trade executes and the mid-price of the best bid and ask quotes at the time of the trade. The minimum possible effective half-spread is the smallest amount by which a liquidity taker could have improved their price by instead providing liquidity. In a standard tick regime, where the minimum price increment is 1¢, this minimum effective half-spread is 0.5¢. In a small tick regime, where the minimum price increment is 0.1¢, it is 0.05¢.

of users with negative PnL into positive territory—roughly one in five losers would have ended the sample with non-negative PnL had they avoided paying the minimum-tick spread.

The pattern reverses at the top of the winner distribution: the spread *increases* elite winners' realized gains, as reflected by negative values of Frac. Spread for the top percentiles. The top 0.1% of winners earn 3.2% *more* than they would absent spread costs—they are net spread *earners*, not net payers. This is consistent with their high maker share documented in Section 4.5 (Table 6), and is the realized-PnL counterpart of the maker–taker wealth transfer we examine in detail in Section 4.7.

Taken together, the decomposition establishes that the bid-ask spread is a meaningful drag on small traders but cannot explain the losses of the worst performers, whose outcomes reflect systematic errors in picking contracts rather than a purely mechanical cost disadvantage.[20]

4.7. Maker-Taker Excess Hit Rate

Section 4.5 showed that users with a higher maker volume share are less likely to lose, and Section 4.6 showed that even a lower-bound estimate of the bid-ask spread accounts for a large share of small-trader losses but a negligible share of the losses of the worst performers. These two results leave open the question of *why* maker volume predicts better performance: is the maker-taker wealth transfer driven mechanically by the liquidity premium captured through the spread, or by superior forecasting skill that liquidity providers bring to their limit-order choices? The excess hit rate, defined as $o_k - p_k$, is constructed from the realized outcome and the traded price itself, so spread

[20] Appendix Tables D7–D11 replicate this decomposition separately for each category. The spread plays its largest role in Politics, where the fraction of losses attributable to transaction costs is 169.2% for the smallest losers (0–20 percentile) and remains above 100% even at the 40–60 percentile (104.1%), implying that many Politics traders would have broken even or profited in the absence of the spread. Culture exhibits a similar pattern, with Frac. Spread reaching 118.9% for the bottom 20% of losers and 59.9% at the 20–40 percentile. Sports, which dominates the platform by user count, has a lower Frac. Spread at the 0–20 percentile (97.9%) than the aggregate (195.8%), consistent with the aggregate figure being pulled up by categories with wider effective spreads such as Politics and Culture. By contrast, Crypto and Tech show markedly lower spread fractions even among the smallest losers—39.1% and 52.3%, respectively.

capture does not mechanically widen it: a maker and a taker who each buy at \$0.40 have the same per-trade excess hit rate. What differs across the two roles is the *selection* of which prices each side transacts at. If maker outperformance were driven entirely by spread capture, the maker-minus-taker excess hit rate gap should be approximately zero at every price bin; a systematic positive gap, and especially one that varies with price, identifies a forecasting-skill component.

Figure 7 plots this gap as a function of the effective buy price, with the complementary sell price on the secondary axis so each bin corresponds to both directions of trade. The solid blue line is the maker-minus-taker excess hit rate; the shaded band with dotted outlines is the 95% confidence interval; the dashed horizontal at zero marks the null of no skill gap; the gray filled area on the right axis shows the fraction of trades at each price; and the vertical line at \$0.96 (equivalently, sell price \$0.04) marks the near-certainty region where the minimum effective spread compresses below one cent and makers can no longer improve on taker prices. A gap that is small and roughly flat across price bins is consistent with the liquidity-premium explanation alone; a gap that is large or rises as prices move away from \$0.5 is harder to reconcile with pure spread capture and points to skill-based selection of when and at what price to post liquidity.

Pooling across categories, the maker-taker excess hit rate gap is positive and tightly estimated at roughly 0.01 across almost the entire range, and collapses to zero past \$0.96 as the spread shrinks—consistent with the liquidity premium being the dominant driver. Cross-sectional variation across categories is large, however. In the most efficient categories, Crypto and Sports, the gap remains small and flat (between 0.01 and 0.02), consistent with liquidity compensation alone. In the other less efficient categories, the gap is both larger (between -0.02 and 0.10) and, critically, rises with deviation from \$0.5: the gap is widest precisely where spread capture would predict no effect, since minimum effective spreads are similar across the mid-range. This price-dependent pattern is the signature of forecasting skill driving maker excess returns on top of the liquidity premium. The one exception runs the other way: in Sports, the gap turns negative at extreme prices (below \$0.1 and above \$0.9). In the microstructure literature, research has shown that when there are large price swings, fast liquidity takers can exploit stale quotes from slower liquidity providers, leading to adverse selection for makers (Baldauf and Mollner, 2022; Grégoire and Martineau, 2022). Sports markets, relative to less actively traded categories such as Politics and Weather, are more likely to experience large price swings even through the last few minutes of trading, offering fast traders more entry points for quick returns (Learner, 2025).

Overall, the evidence in this section complements Sections 4.5 and 4.6: maker volume predicts lower loss probability and the spread explains small-trader losses but not top-performer profits; the maker-taker excess hit rate gap shows that, on top of the mechanical liquidity premium, makers earn a genuine forecasting premium that is concentrated in the categories where prices are least efficient and mispricings are hardest for takers to exploit.

5. Conclusion and implications for future research

Our analysis of over 2.4 million Polymarket users reveals five findings. Profits are concentrated: the top 1% of users capture 76.5% of all gains, while the median user loses money, consistent with a near-zero-sum market. Prices are relatively efficient and informative. A contract priced at p resolves in its favor approximately p percent of the time, so the market as a whole aggregates information effectively. Yet this efficiency is itself the central hazard for investors: because prices already reflect the collective wisdom of all traders, systematic outperformance requires identifying genuine mispricings, a skill concentrated in a tiny number of sophisticated traders. Investors entering prediction markets should therefore treat them as they would other highly competitive financial markets: on average, prices are fair, the gross expected return is zero by construction, and the expected return is negative for users who take liquidity, conditional on the spread and information environment they face.

This paper is a first step of a broader research agenda on investor behavior in prediction markets. Several questions remain open. First, while we identify extreme-price trading and low market-making activity as predictors of loss, the deeper behavioral mechanisms such as overconfidence, probability misperception, or simply a preference for lottery-like payoffs deserve more empirical investigation using trader-level data. Second, as prediction markets grow in scale and visibility, understanding their interaction with traditional financial markets becomes increasingly important: Polymarket prices on macroeconomic and political events are now widely cited by equity analysts and news media, raising questions about whether they independently move asset prices or merely reflect information already embedded in them. Prediction markets carry genuine informational content, yet the traders, their behavior, and how it leads to price discovery remain poorly understood. Third, our sample ends on the day before Polymarket introduced taker fees on all markets. Because fees materially change the economics of liquidity-taking and high-frequency market-making, the cross-section of winners and losers may look different in the post-fee regime; assessing the external validity of our findings under that new fee structure is an important question we leave for future research. We intend to pursue these questions in future work.

References

Ayers, John W., Atharva Yeola, Matthew R. Allen, Nimit Desai, Adam Poliak, Kevin H. Yang, and Davey M. Smith, 2025, Growing health concern regarding gambling addiction in the age of sportsbooks, *JAMA Internal Medicine* 185, 382–389.

Baker, Scott R., Justin Balthrop, Mark J. Johnson, Jason D. Kotter, and Kevin Pisciotta, 2024, Gambling away stability: Sports betting's impact on vulnerable households, .

Baldauf, Markus, and Joshua Mollner, 2022, Fast traders make a quick buck: The role of speed in liquidity provision, *Journal of Financial Markets* 58, 100621.

Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, 2009, Just how much do individual investors lose by trading?, *Review of Financial Studies* 22, 609–632.

Barber, Brad M, and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *The journal of Finance* 55, 773–806.

Becker, Jonah, 2025, The microstructure of wealth transfer in prediction markets, `https://www.jbecker.dev/research/prediction-market-microstructure` Accessed: 2026-03-15.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2019, Price discovery without trading: Evidence from limit orders, *The Journal of Finance* 74, 1621–1658.

Bryzgalova, Svetlana, Anna Pavlova, and Taisiya Sikorskaya, 2023, Retail trading in options and the rise of the big three wholesalers, *Journal of Finance* 78, 3465–3514.

Cookson, J. Anthony, 2018, When saving is gambling, *Journal of Financial Economics* 128, 24–50.

Cziraki, Peter, and Jasmin Gider, 2021, The dollar profits to insider trading, *Review of Finance* 25, 1547–1580.

Diercks, Anthony M, Jared Dean Katz, and Jonathan H Wright, 2026, Kalshi and the rise of macro markets, *Working paper*.

Forsythe, Robert, Forrest Nelson, George R. Neumann, and Jack Wright, 1992, Anatomy of an experimental political stock market, *American Economic Review* 82, 1142–1161.

Gomez Cram, Roberto, Yunhan Guo, Theis Ingerslev Jensen, and Howard Kung, 2025, Financial prediction markets: A new measure of earnings expectations, *Available at SSRN 5933475*.

Grégoire, Vincent, and Charles Martineau, 2022, How is earnings news transmitted to stock prices?, *Journal of Accounting Research* 60, 261–297.

Hayek, Friedrich A., 1945, The use of knowledge in society, *American Economic Review* 35, 519–530.

Kaniel, Ron, and Hong Liu, 2006, So what orders do informed traders use?, *The Journal of Business* 79, 1867–1913.

Learner, Sam, 2025, Prediction markets barely make money; sportsbooks make money, *Financial Times* Accessed: 2026-04-17.

Manski, Charles F., 2006, Interpreting the predictions of prediction markets, *Economics Letters* 91, 425–429.

Newall, Philip, 2025, Sludge, dark patterns and dark nudges: A taxonomy of online gambling platforms' deceptive design features, *Addiction* 120, 1916–1923.

Oliven, Kenneth, and Thomas A. Rietz, 2004, Suckers are born but markets are made: Individual rationality, arbitrage and market efficiency on an electronic futures market, *Management Science* 50, 336–351.

Packin, Nizan Geslevich, and Sharon Rabinovitz, 2026, Prediction markets as a public health threat, *Science* 392, 257–260.

Reichenbach, Felix, and Miriam Walther, 2025, Exploring decentralized prediction markets: Accuracy, skill, and bias on Polymarket, *Working paper.*

Sirolly, Arman, Haochen Ma, Yash Kanoria, and Rajiv Sethi, 2025, Network-based detection of wash trading in decentralized prediction markets, *Working paper.*

Snowberg, Erik, and Justin Wolfers, 2010, Explaining the favorite–long shot bias: Is it risk-love or misperceptions?, *Journal of Political Economy* 118, 723–746.

——— , and Eric Zitzewitz, 2012, Prediction markets for economic forecasting, *Handbook of Economic Forecasting* 2, 827–880.

Thaler, Richard H., and William T. Ziemba, 1988, Anomalies: Parimutuel betting markets: Racetracks and lotteries, *Journal of Economic Perspectives* 2, 161–174.

Tsang, Kwok Ping, and Zichao Yang, 2025, The anatomy of Polymarket: Evidence from the 2024 presidential election, *Working Paper.*

Wolfers, Justin, and Eric Zitzewitz, 2004, Prediction markets, *Journal of Economic Perspectives* 18, 107–126.

Figure 1: Polymarket User and Trading Volume Growth Over Time

This figure shows the growth of Polymarket over time. Panel A shows the number of new users per month and the cumulative number of users over time. Panel B shows the total trading volume per month (in Billion USD) per category and the cumulative trading volume over time. The sample period is from November 11, 2022 to March 29, 2026.

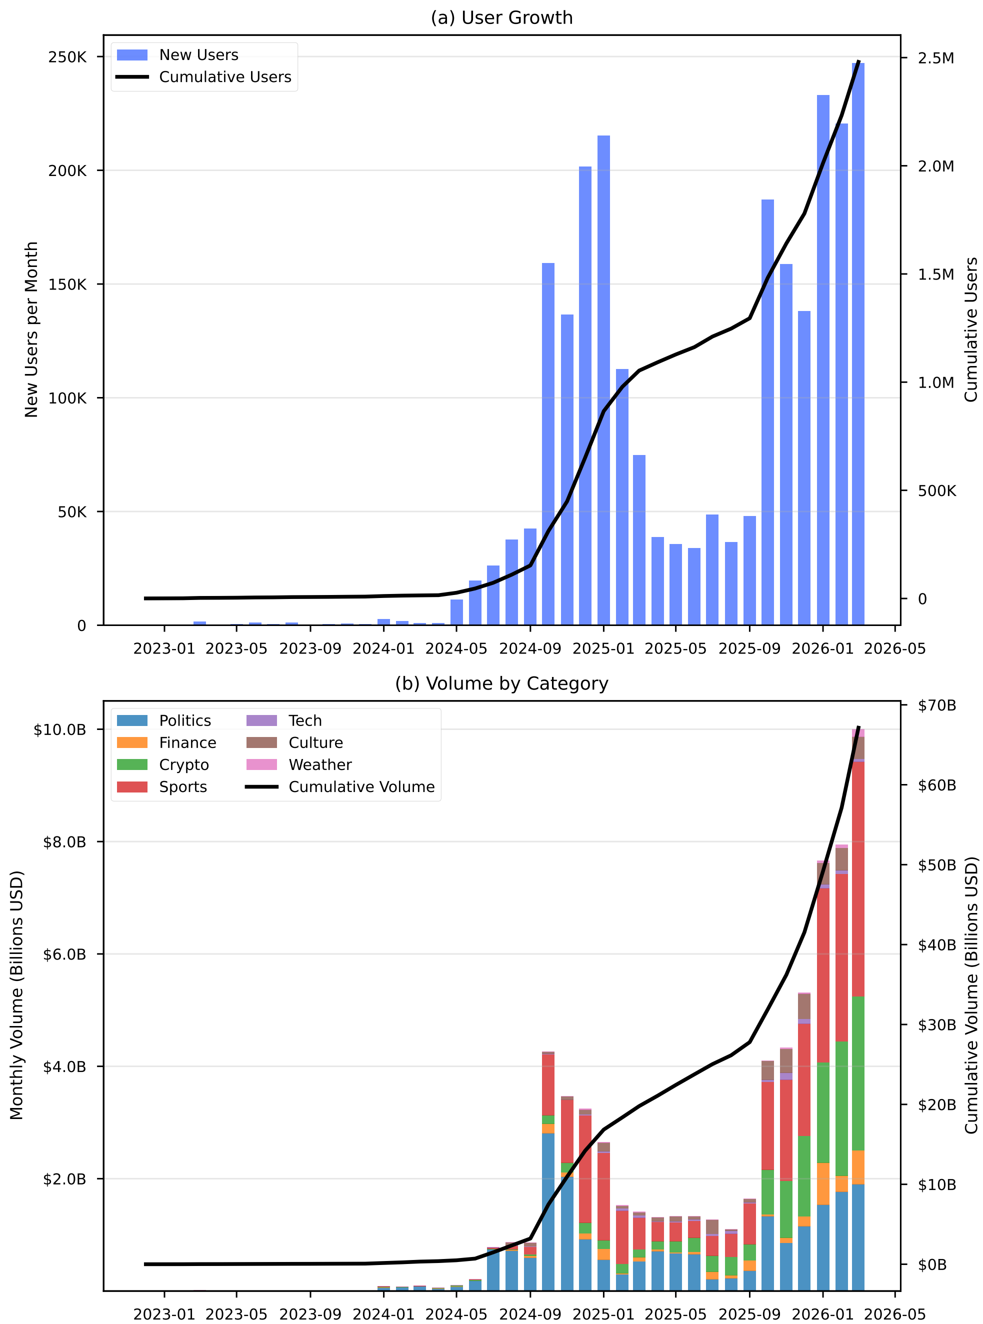

Figure 2: Polymarket's Number of New Markets

This figure shows the number of new events and markets created on Polymarket over time. The sample period is from November 11, 2022 to March 29, 2026.

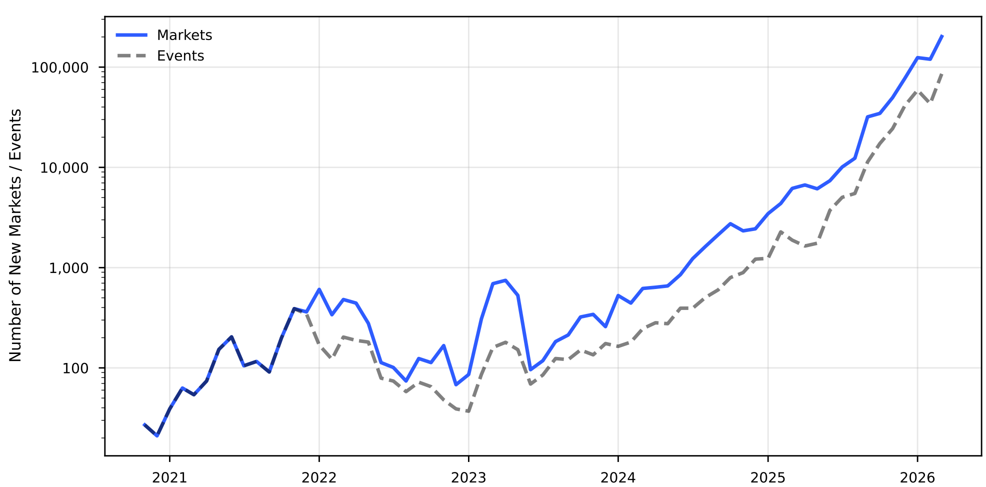

Figure 3: Concentration of Gains and Losses Across Polymarket Users

This figure shows the cumulative concentration of trading gains and losses among Polymarket users. For each percentile threshold (Top 0.1%, 1%, 5%, 10%, 25%, 50%), the green bar reports the share of total gains captured by users in that top group (winners ranked from highest to lowest PnL), and the red bar reports the share of total losses borne by users in that bottom group (losers ranked from largest to smallest loss). The sample period is from November 11, 2022 to March 29, 2026.

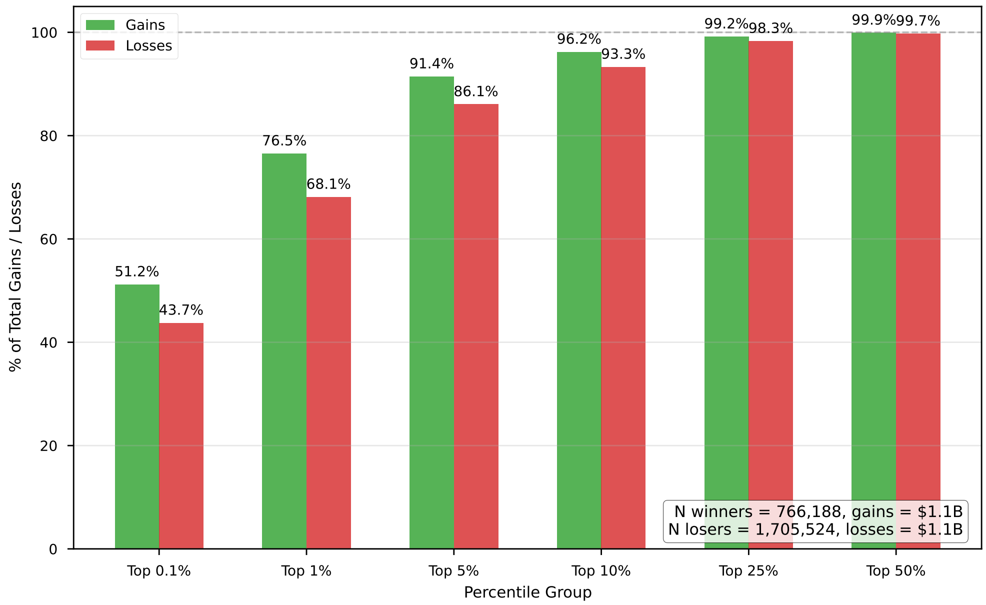

Figure 4: Concentration of Gains and Losses Across Polymarket Users

Panel A of this figure shows the concentration of gains and losses across Polymarket users over time for the top 0.1%, 1%, and 5% gains and losses. Panel B shows the fraction of users with positive PnL over time.

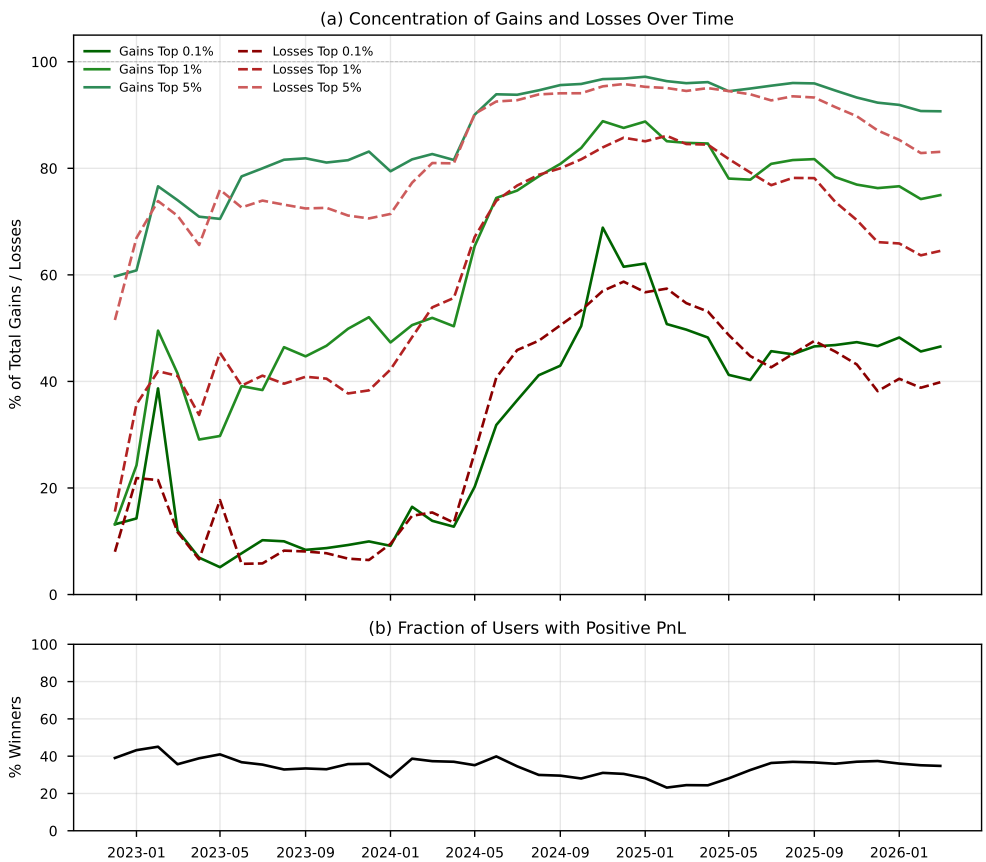

Figure 5: Calibration of Polymarket Prices

Panel A of this figure plots the calibration of Polymarket binary market prices at different time horizons before resolution. The remaining panels show the same calibration broken down by market category. Market prices are binned into ten equal-width bins (0–10¢, 10–20¢, . . . , 90¢–$1). For each bin, the x-axis shows the midpoint of the bin (the expected probability) and the y-axis shows the observed frequency of the "Yes" outcome among markets in that bin (the actual probability). The dashed 45-degree line represents perfect calibration. Shaded bands show 95% bootstrap confidence intervals for the observed frequency (1,000 replications). Only markets with valid prices at all displayed horizons are included, and categories with fewer than ten such markets are dropped from the figure. The sample period is from November 11, 2022 to March 29, 2026.

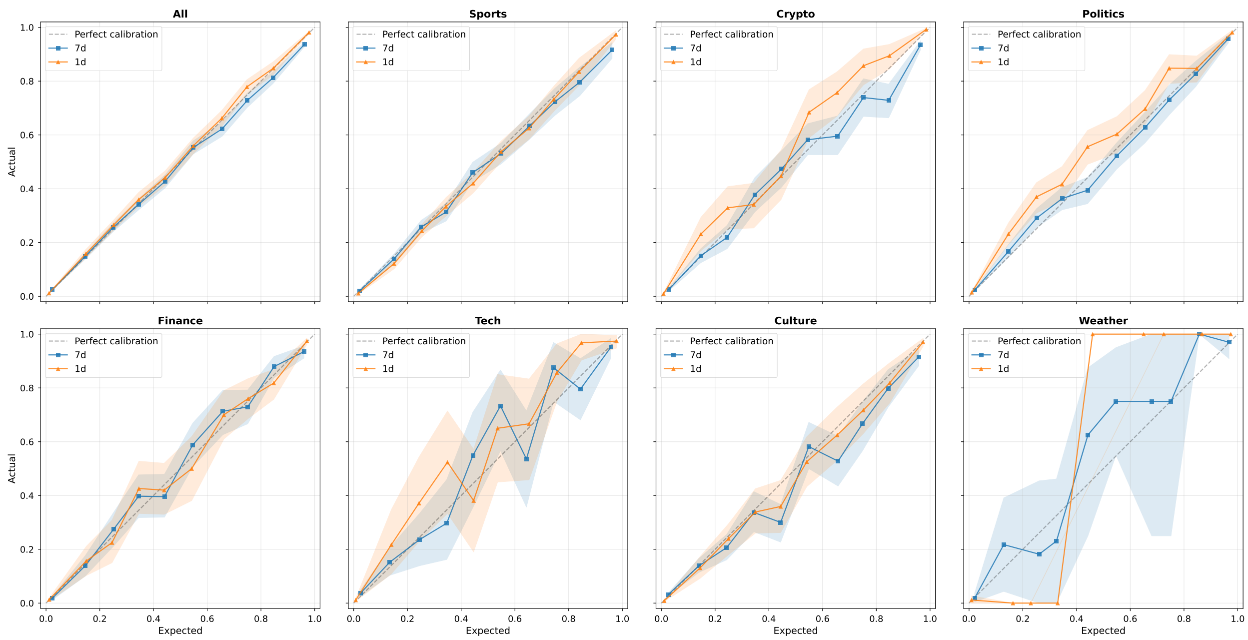

Figure 6: Prediction Accuracy of Winners and Losers

This figure plots the mean excess hit rate $\text{EHR}_i = \frac{1}{K_i}\sum_k (o_k - p_k)$ (top panel) and the volume-weighted excess hit rate $\text{EHR}_{\text{vw},i} = \sum_k q_k(o_k - p_k)/\sum_k q_k$ (bottom panel) by PnL percentile, ranking users from most profitable (left) to least profitable (right). Users are sorted by PnL and assigned to 200 half-percentile bins. Within each bin we plot the mean (solid blue, with a 95% confidence band) and the median (dashed black). Vertical dotted lines mark the bins that contain the PnL thresholds of \$100, \$0, and $-$\$100. Each trade k is normalized to the buy-side perspective: a sell of outcome X at price p is recast as a buy of the complementary contract at $1 - p$. $o_k \in \{0, 1\}$ equals 1 iff the contract purchased in trade k (after buy-side normalization) resolves in-the-money, p_k is the buy-side normalized trade price, and q_k is the traded quantity in shares (equivalently, the trade's USDC notional under the matched-notional convention). A positive value indicates that traded contracts resolve more favorably than their prices implied; a negative value indicates the opposite. The sample period is from November 11, 2022 to March 29, 2026.

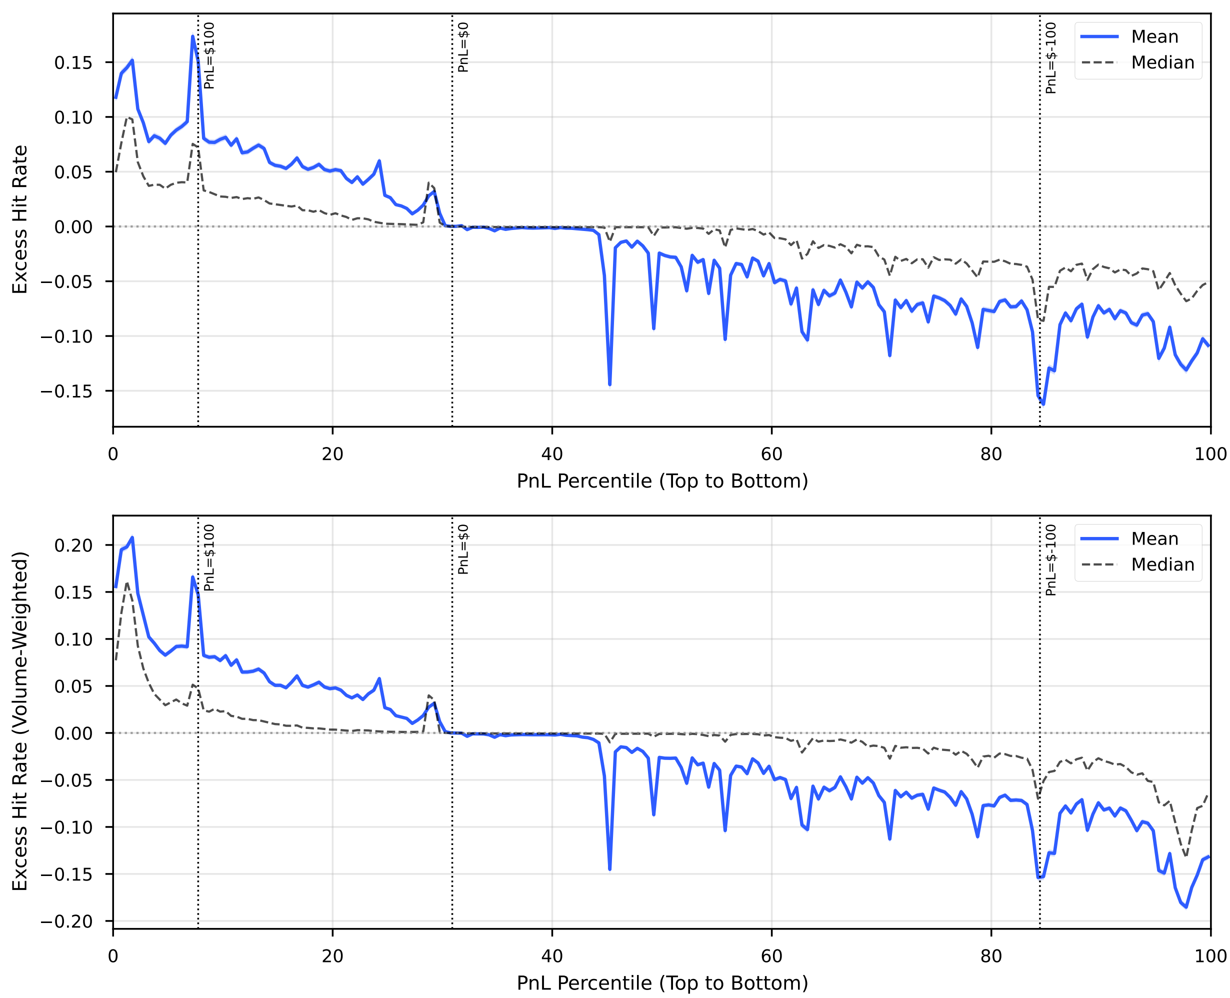

Figure 7: Excess Hit Rate Differences Between Makers and Takers

This figure plots the difference in excess hit rate between makers and takers across price levels, separately for all categories combined and for each individual category. The solid blue line shows the maker-minus-taker difference in excess hit rate at each price level, with shaded bands bounded by dotted outlines representing 95% confidence intervals. The dashed horizontal line at zero marks the null of no maker-taker gap. The gray filled area plots the fraction of total trades at each price bin on the right y-axis. The x-axis is split into two regions: the effective buy price (ranging from \$0.50 to \$1.00) and the effective sell price (ranging from \$0.01 to \$0.50). The vertical line marks the effective spread boundary at \$0.96 for buys and \$0.04 for sells, beyond which maker quotes are required to cross the spread and quote subpenny prices. The sample period is from November 11, 2022 to March 29, 2026.

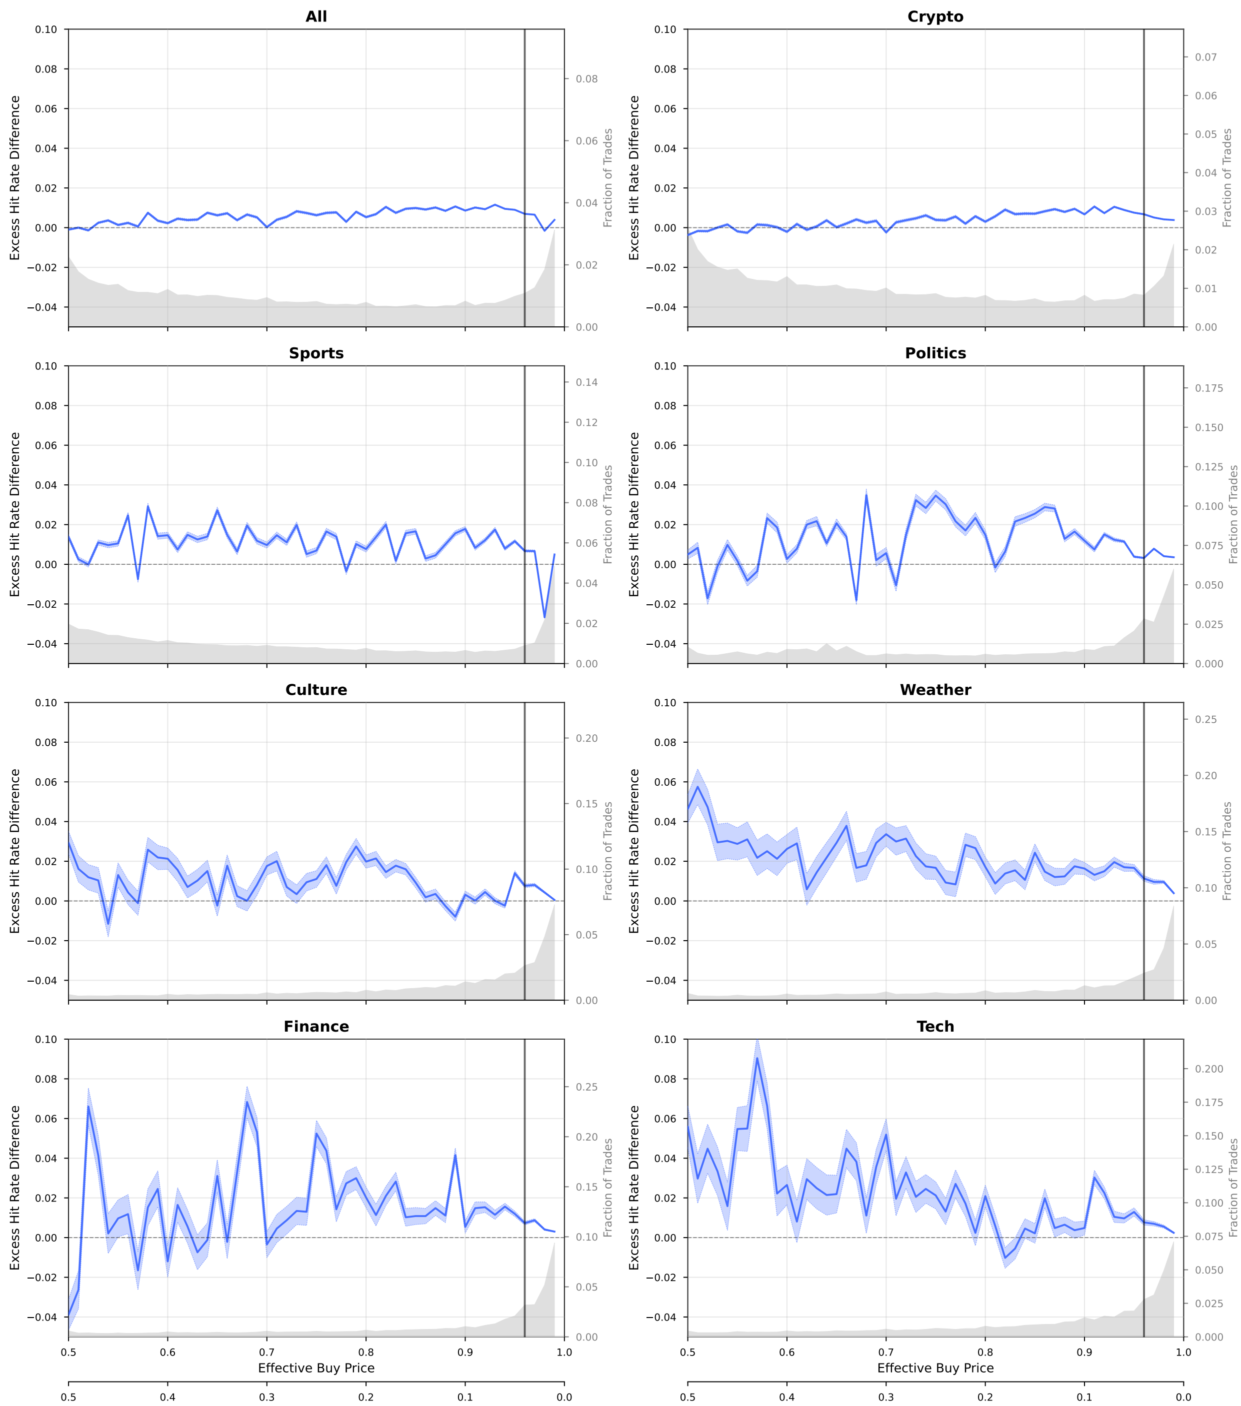

Table 1: Distribution of Trading Activity Across Polymarket Users

This table reports user-level summary statistics for four metrics across all Polymarket users in our sample. *Trades per User* is the total number of executed trades. *Markets per User* is the number of distinct prediction market questions traded. *PnL per User* is the final mark-to-market profit and loss in USD, where the USDC balance captures realized gains and losses and the portfolio value captures unrealized positions. *Fraction Maker Volume* is the share of a user's traded volume that they supplied as a passive limit order (maker). Statistics reported are the minimum, mean, 5th, 25th, 50th, 75th, and 95th percentiles, and maximum across all users. The sample period is from November 11, 2022 to March 29, 2026.

Metric	Min	Mean	5th Pct.	25th Pct.	Median	75th Pct.	95th Pct.	Max
Trades per User	1	474	1	9	28	74	546	13,587,640
Questions per User	1	54	1	3	10	26	127	130,387
PnL per User (USD)	-$10,022,403	$0	-$920	-$32	-$2	$1	$252	$22,029,973
Fraction Maker Volume	0.0%	17.9%	0.0%	0.0%	0.0%	39.5%	75.4%	100.0%

Table 2: Summary Statistics by Category

This table reports trading activity broken down by market category. *Events* and *Markets* are the number of distinct events and markets. *Trades* and *Volume* are reported in levels and as a percentage of total. *Users* is the number of unique wallet addresses appearing as maker or taker in that category; percentages are relative to the total number of unique users across all categories (users trading in multiple categories are counted once in the total). *Avg Trades/M* and *Avg Vol/M* are per-events and per-markets averages. The event and market counts reported here are restricted to the universe that enters the trade-level analyses, i.e., those with at least one reconciled trade in the sample, and are therefore smaller than the classification universe of 316,429 events and 729,133 markets described in Section A.2. The sample period is from November 11, 2022 to March 29, 2026.

Category	Events	Markets	Trades (M)	Trades (%)	Volume ($M)	Volume (%)	Users	Users (%)	Avg Trades/M	Avg Vol/M
Sports	64,205	321,231	95.4	16.2	25,477.0	37.9	1,573,302	63.4	296.9	$79,311
Politics	7,248	29,926	55.3	9.4	21,114.4	31.5	1,568,028	63.2	1,847.6	$705,553
Crypto	172,313	206,491	394.7	67.1	12,762.0	19.0	1,373,525	55.4	1,911.4	$61,804
Culture	2,530	15,806	19.5	3.3	3,473.9	5.2	732,477	29.5	1,235.0	$219,781
Finance	5,539	17,159	9.2	1.6	3,107.5	4.6	720,184	29.0	534.7	$181,100
Tech	1,016	4,308	5.9	1.0	806.9	1.2	469,579	18.9	1,363.6	$187,303
Weather	2,576	19,962	8.3	1.4	393.7	0.6	199,565	8.0	417.7	$19,725
Total	255,427	614,883	588.3	100.0	67,135.4	100.0	2,480,080	100.0	956.7	$109,184

Table 3: PnL Concentration by Category

This table reports descriptive statistics and the distribution of mean PnL across percentile buckets, broken down by market category. The top panel reports the number of markets, events, users, trades, and total volume for each category, as well as the fraction of users with positive PnL. The middle and bottom panels report mean PnL by percentile bucket for losers (PnL $<$ 0) and winners (PnL $>$ 0), respectively. Percentile buckets are constructed within each category independently. The sample period is from November 11, 2022 to March 29, 2026.

	All	Sports	Politics	Crypto	Culture	Finance	Tech	Weather
Descriptive Statistics								
N Markets	614,884	321,231	29,926	206,491	15,806	17,159	4,308	19,962
N Events	255,427	64,205	7,248	172,313	2,530	5,539	1,016	2,576
N Users (000s)	2,480.1	1,573.3	1,568.0	1,373.5	732.5	720.2	469.6	199.6
N Trades (M)	588.3	95.4	55.3	394.7	19.5	9.2	5.9	8.3
Volume (000s)	$67,140,742	$25,477,039	$21,114,381	$12,762,004	$3,473,853	$3,107,499	$806,902	$393,742
Frac. Winners	30.9%	31.5%	35.5%	42.5%	38.6%	46.6%	50.3%	52.6%
Mean PnL by Percentile – Losers (PnL $<$ 0)								
0–20	$-0.37	$-0.22	$-0.22	$-1.07	$-0.07	$-0.10	$-0.08	$-0.39
20–40	$-3.23	$-1.88	$-2.09	$-10	$-0.85	$-1.06	$-1.07	$-3.46
40–60	$-14	$-10	$-9.32	$-41	$-5.26	$-6.61	$-6.57	$-14
60–80	$-61	$-56	$-53	$-183	$-30	$-33	$-31	$-48
80–90	$-256	$-251	$-275	$-757	$-132	$-146	$-118	$-153
90–95	$-892	$-863	$-1,029	$-2,030	$-479	$-530	$-385	$-426
95–99	$-2,760	$-3,134	$-4,886	$-6,175	$-2,350	$-2,467	$-1,663	$-1,613
99–99.9	$-16,675	$-20,144	$-37,090	$-27,863	$-19,200	$-18,861	$-10,663	$-8,117
99.9–100	$-269,401	$-327,714	$-441,611	$-177,034	$-263,777	$-199,154	$-98,724	$-52,959
Mean PnL by Percentile – Winners (PnL $>$ 0)								
0–20	$0.18	$0.08	$0.07	$0.10	$0.03	$0.02	$0.02	$0.01
20–40	$2.14	$2.30	$0.88	$1.15	$0.42	$0.24	$0.25	$0.16
40–60	$13	$13	$4.94	$7.08	$2.97	$1.89	$1.52	$1.29
60–80	$72	$61	$27	$34	$15	$12	$9.09	$9.46
80–90	$345	$293	$121	$166	$55	$51	$33	$34
90–95	$1,306	$1,108	$444	$686	$162	$163	$98	$98
95–99	$5,108	$4,046	$2,452	$4,146	$881	$882	$499	$581
99–99.9	$38,599	$30,626	$23,752	$23,931	$8,421	$7,512	$5,021	$4,050
99.9–100	$701,648	$595,603	$629,959	$217,583	$164,905	$171,620	$56,911	$50,548

Table 4: Transition Matrix of PnL Groups

This table reports transition probabilities across monthly-performance groups over 1-month, 3-month, and 6-month horizons, providing a test of whether strong or weak monthly performance persists (as opposed to the persistence of cumulative PnL, which is mechanically strong). In each month, users who traded are sorted into five groups based on the change in their mark-to-market profit and loss (PnL) during that month, i.e., the PnL accrued or realized over the month, with thresholds < -\$100, -\$100 to -\$10, -\$10 to \$10, \$10 to \$100, and > \$100. The table shows the percentage of users in each starting group whose monthly PnL change falls into each group after the specified horizon, as well as the percentage that drop out (i.e., no trades in the target month). The sample period is from November 11, 2022 to March 29, 2026.

Panel A: 1-Month Horizon							
Starting Group	N	< -\$100	-\$100 to -\$10	-\$10 to \$10	\$10 to \$100	> \$100	Dropped
< -\$100	275,640	20.8%	9.1%	8.8%	5.2%	12.0%	44.1%
-\$100 to -\$10	437,629	4.6%	14.0%	26.8%	8.8%	3.0%	42.7%
-\$10 to \$10	3,260,841	0.7%	3.3%	56.9%	2.4%	0.5%	36.2%
\$10 to \$100	307,046	6.6%	15.8%	26.9%	15.0%	4.8%	31.0%
> \$100	193,705	23.7%	6.9%	9.5%	6.7%	24.6%	28.6%
Panel B: 3-Month Horizon							
Starting Group	N	< -\$100	-\$100 to -\$10	-\$10 to \$10	\$10 to \$100	> \$100	Dropped
< -\$100	187,506	14.0%	6.3%	9.2%	4.6%	9.3%	56.6%
-\$100 to -\$10	314,294	3.4%	8.4%	22.4%	5.7%	2.3%	57.7%
-\$10 to \$10	2,768,522	0.7%	2.3%	42.2%	1.8%	0.5%	52.5%
\$10 to \$100	231,008	4.0%	8.7%	25.9%	8.9%	3.3%	49.3%
> \$100	139,848	15.5%	6.0%	10.3%	5.4%	18.3%	44.5%
Panel C: 6-Month Horizon							
Starting Group	N	< -\$100	-\$100 to -\$10	-\$10 to \$10	\$10 to \$100	> \$100	Dropped
< -\$100	110,155	11.1%	4.6%	6.8%	3.3%	8.2%	66.1%
-\$100 to -\$10	206,858	2.6%	5.3%	17.3%	3.6%	1.9%	69.2%
-\$10 to \$10	2,097,292	1.1%	1.9%	31.4%	1.5%	0.6%	63.5%
\$10 to \$100	153,864	3.0%	5.3%	19.5%	5.3%	2.8%	64.1%
> \$100	83,826	12.6%	4.6%	8.0%	4.1%	15.8%	55.0%

Table 5: Probit Regression: Determinants of Losses

This table reports probit regressions where the dependent variable is an indicator equal to one if the user's PnL is negative. Reported values are marginal effects at the sample mean, computed from the fitted probit as $\phi(\bar{x}'\hat{\beta})\hat{\beta}_k$. The independent variables include the fraction of trades at extreme prices (below 10¢ or above 90¢), the fraction of maker volume, log number of trades, log total volume, the category HHI (Herfindahl-Hirschman Index of volume concentration across market categories) and counterparty HHI as a proxy for wash trading. Column (4) further adds a set of category fixed effects: a binary indicator for each of the seven market categories equal to one if the user has any trading volume in that category, with Sports omitted as the reference. Robust standard errors are reported in parentheses. The table reports the results for all users (columns 1–4), for users with more than 100 trades (columns 5–6), and for users with more than 1,000 trades (columns 7–8). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

	All users				Users with more than 100 trades		Users with more than 1,000 trades	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Frac Extreme Price	-0.012^{***}	-0.016^{***}	-0.006^{***}	-0.003^{***}	0.007^{***}	0.012^{***}	-0.129^{***}	-0.128^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.006)	(0.006)
Frac Maker Volume		-0.354^{***}	-0.349^{***}	-0.344^{***}	-0.327^{***}	-0.324^{***}	-0.147^{***}	-0.142^{***}
		(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.006)	(0.006)
Log N Trades		-0.001^{***}	0.004^{***}	0.017^{***}	0.027^{***}	0.029^{***}	-0.004^{**}	-0.006^{***}
		(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.002)	(0.002)
Log Total Volume		0.020^{***}	0.018^{***}	0.020^{***}	-0.021^{***}	-0.021^{***}	-0.043^{***}	-0.041^{***}
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Category HHI			0.081^{***}	-0.055^{***}	-0.010^{***}	-0.067^{***}	-0.131^{***}	-0.143^{***}
			(0.001)	(0.002)	(0.003)	(0.004)	(0.007)	(0.010)
Counterparty HHI			-0.005^{***}	0.036^{***}	0.180^{***}	0.206^{***}	0.500^{***}	0.494^{***}
			(0.002)	(0.002)	(0.007)	(0.007)	(0.050)	(0.049)
Traded Crypto				-0.074^{***}		-0.016^{***}		0.057^{***}
				(0.001)		(0.002)		(0.006)
Traded Finance				-0.033^{***}		-0.018^{***}		-0.010^{*}
				(0.001)		(0.002)		(0.005)
Traded Politics				-0.056^{***}		0.000		0.041^{***}
				(0.001)		(0.002)		(0.005)
Traded Tech				-0.033^{***}		-0.023^{***}		-0.026^{***}
				(0.001)		(0.002)		(0.005)
Traded Culture				0.002^{***}		-0.006^{***}		-0.024^{***}
				(0.001)		(0.002)		(0.005)
Traded Weather				-0.036^{***}		-0.021^{***}		-0.008
				(0.001)		(0.002)		(0.005)
Observations	2,480,101	2,480,101	2,480,101	2,480,101	478,095	478,095	78,366	78,366
Pseudo R^2	0.000	0.034	0.036	0.043	0.041	0.042	0.054	0.056

Table 6: PnL Group Characteristics

This table reports trading characteristics for users sorted into four non-overlapping groups based on their total mark-to-market profit and loss (PnL). The "Top 0.1%" column reports means for the 0.1% of users with the highest PnL; "Next 0.9%" for the next 0.9% of users (the remainder of the top 1%, excluding the top 0.1%); "Next 4%" for the users in the top 5% not already in the top 1%; and "Bot. 95%" for the bottom 95% of users. *Excess Hit Rate* is the user's average of $(o_k - p_k)$ across trades, where $o_k \in \{0, 1\}$ equals 1 iff the contract purchased in trade k (after buy-side normalization) resolves in-the-money and p_k is the buy-side normalized trade price; see equation (1). *Frac Maker Volume* is the fraction of a user's traded volume supplied as a passive limit order (maker). *Frac Extreme Price* is the fraction of a user's trades executed at prices below 10¢ or above 90¢. *Category HHI* is the Herfindahl-Hirschman Index of a user's trading volume concentration across market categories, ranging from 0 (fully diversified) to 1 (all volume in one category). Columns (5)–(7) report the difference in means between each top group and the bottom 95%, with significance levels from Welch's two-sample t-test: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The sample period is from November 11, 2022 to March 29, 2026.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	PnL group mean				Difference		
	Top 0.1%	Next 0.9%	Next 4%	Bot. 95%	0.1%–95%	0.9%–95%	4%–95%
Excess Hit Rate	0.0974	0.1317	0.1027	-0.0265	0.1238***	0.1582***	0.1292***
Frac Maker Volume	0.4721	0.4140	0.3236	0.1706	0.3016***	0.2434***	0.1530***
Frac Extreme Price	0.2820	0.2870	0.3115	0.5642	-0.2822***	-0.2772***	-0.2527***
Category HHI	0.8038	0.8165	0.7949	0.7670	0.0368***	0.0495***	0.0279***

Table 7: PnL Spread Decomposition Statistics

This table reports mean profit and loss (PnL) including and excluding a lower-bound estimate of spread costs, by percentile bucket, separately for users with negative PnL (Panel A: Losers) and positive PnL (Panel B: Winners). Users within each group are ranked by the magnitude of their PnL and assigned to percentile buckets: 0–20, 20–40, 40–60, 60–80, 80–90, 90–95, 95–99, 99–99.9, and 99.9–100. *Mean PnL (incl.)* is the average mark-to-market PnL including spread costs. *Mean PnL (excl.)* is the counterfactual PnL after removing the minimum possible effective half-spread on each taker trade: 0.5¢ in the standard tick regime and 0.05¢ in the small-tick regime (prices below 4¢ or above 96¢). Because the minimum tick is a floor on the effective half-spread rather than its typical value, this adjustment is a lower bound on spread-related costs and, correspondingly, an upper bound on the share of losses that cannot be attributed to the spread. *Difference* is the gap between the two, tested for statistical significance using a t-test. *Frac. Spread* is the fraction of the average gain or loss attributable to the spread, computed as $(\text{PnL}_{\text{excl.}} - \text{PnL}_{\text{incl.}})/|\text{PnL}_{\text{incl.}}|$. *Maker %* is the average fraction of volume supplied as a passive limit order (maker) within each bucket. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL $<$ 0)						
0–20	341,105	-0.37	0.36	-0.73^{***}	195.8%	13.1%
20–40	341,105	-3.23	0.36	-3.60^{***}	111.1%	11.0%
40–60	341,105	-13.58	-5.45	-8.13^{***}	59.8%	10.1%
60–80	341,105	-60.87	-50.65	-10.22^{***}	16.8%	14.0%
80–90	170,552	-256.41	-235.79	-20.62^{***}	8.0%	17.9%
90–95	85,276	-892.34	-859.76	-32.58^{***}	3.7%	28.6%
95–99	68,221	-2,760.45	-2,690.62	-69.83^{***}	2.5%	35.4%
99–99.9	15,350	-16,675.11	-16,359.73	-315.38^{***}	1.9%	29.3%
99.9–100	1,705	-269,401.48	-266,976.23	$-2,425.25$	0.9%	37.3%
Panel B: Winners (PnL $>$ 0)						
0–20	153,238	0.18	0.47	-0.29^{***}	160.3%	24.4%
20–40	153,238	2.14	3.12	-0.98^{***}	45.6%	23.1%
40–60	153,237	13.27	15.14	-1.87^{***}	14.1%	20.6%
60–80	153,238	72.33	75.05	-2.73^{***}	3.8%	24.9%
80–90	76,619	344.79	352.91	-8.12^{***}	2.4%	26.3%
90–95	38,309	1,306.33	1,302.12	4.21	-0.3%	35.1%
95–99	30,648	5,108.26	5,044.72	63.54^{***}	-1.2%	40.6%
99–99.9	6,895	38,598.91	37,160.71	$1,438.21^{***}$	-3.7%	44.8%
99.9–100	766	701,648.09	678,885.67	$22,762.42^{***}$	-3.2%	49.3%

Appendix Material for

Who Wins and Who Loses In Prediction Markets?

Pat Akey	Vincent Grégoire	Nicolas Harvie	Charles Martineau
ESSEC Business School	HEC Montréal	University of Toronto	University of Toronto

Appendix

A. Data Collection

We assemble our dataset by combining the outputs of four public Polymarket APIs.

Gamma API. Polymarket's Gamma REST API returns JSON records describing markets and events. For each market, we retrieve the question text, slug, description, condition identifier, outcome tokens (token IDs and labels), cumulative volume and liquidity, open and close timestamps, resolution status, the winning outcome for resolved markets, and the platform-assigned tags that we use to classify markets into the seven categories (Sports, Crypto, Politics, Finance, Tech, Culture, Weather). Events group related markets and provide the parent title, description, and tags.

SubGraph API. Trade-level data come from the Polymarket SubGraph, a Goldsky-hosted GraphQL index of on-chain order-filled events from the Polygon blockchain. Each event returns the transaction hash, block timestamp, maker and taker wallet addresses, the maker and taker asset (token) identifiers, the filled amounts in each direction, and the fee paid. Because the SubGraph indexes every on-chain settlement, this yields the complete population of settled trades.[21]

CLOB API. Polymarket's Central Limit Order Book (CLOB) REST API exposes market-level data used by the trading interface. We rely on it specifically for the "simplified markets" dataset, a compact listing of each market's outcome tokens together with an explicit Boolean flag indicating which token, if any, has been declared the winner by the resolution oracle. This flag is the primary signal behind the outcome-attribution procedure described in Section A.3.

Data API. Polymarket's Data REST API exposes user-level information keyed by wallet address. We use it to retrieve user positions and PnL, which we use to validate our reconstructed PnL series

[21] Polymarket's Data API also exposes a per-market trades endpoint, which in principle provides similar information. However, the number of trades that can be retrieved per market is capped, which limits visibility on the most active markets. The SubGraph has no such cap.

(see Section 3.2).

The trade-level SubGraph records are then reconciled with the Gamma market metadata, so that every trade is linked to its underlying question, outcome tokens, and final resolution.

A.1. Reconciling Trades from Order-Filled Events

All Polymarket trades settle through a single smart contract, the CTF (Conditional Token Framework) Exchange (or its negative-risk variant for multi-outcome markets). The SubGraph exposes the raw stream of on-chain *OrderFilled* events emitted by this contract, each of which records a maker address, a taker address, the asset identifiers and amounts on each side of the fill, and the transaction hash. Because the exchange contract is an infrastructure layer that holds custody of assets during settlement rather than an economic counterparty, it appears in the maker or taker field of most events. Taking the raw events at face value would therefore both misidentify the counterparty and inflate volume. The purpose of the reconciliation step is to recover, from this raw event stream, the two actual end-user wallets involved in each trade and to label which of them consumed liquidity (taker) and which provided it (maker).

The reconciliation relies on a fundamental distinction between two types of fills. A *direct fill* is a single-event transaction in which a user's resting order is matched against an off-exchange operator. In that case, the maker is the end user and the taker is the operator; both addresses already correspond to real economic agents, and the event is retained as-is. A *matched trade*, which is by far the more common case, occurs when the exchange internally crosses a taker order against one or more resting maker orders within a single on-chain transaction. In that case, the contract emits multiple OrderFilled events sharing the same transaction hash: one event per maker order, plus one additional "taker-order" event in which the maker field records the taker user and the taker field records the exchange contract itself.

The key insight is that, within a matched-trade transaction, the identity of the end-user taker can be recovered from the maker field of the single event whose taker field equals the exchange address. Once that address is known, the other events in the same transaction (the maker orders) already carry the correct maker (the liquidity-providing user) and need only have their taker field rewritten from the exchange contract to the recovered taker user. The taker-order event is then discarded, because keeping both sides of the cross would double-count the trade.

Operationally, we flag each event according to whether its taker is one of the two exchange contracts, group events by transaction hash, and derive for each transaction a Boolean indicator of whether it is a matched trade as well as, when applicable, the recovered taker-user address. We then drop the taker-order events from matched trades, rewrite the taker field of the remaining events with the recovered address, and leave direct fills untouched. When a single taker order crosses against N distinct makers, this procedure yields N rows, one per maker–taker pair, each carrying the fraction of the taker's order filled against that particular maker. Volume is therefore conserved and each reconciled row corresponds to a unique economic trade between two end users, with the maker-versus-taker designation preserved from the underlying limit-order-book interaction. A final set of validation checks verifies that no exchange address remains in either the maker or taker field of the reconciled dataset.

A.2. Assigning Markets to Categories

Polymarket does not expose a clean, market-level category field. The platform attaches a free-form list of "tags" to each event (the container that groups related markets under a single question) and separately stores a coarse legacy "category" string on the event record. Individual markets inherit their classification from the parent event; markets that are not linked to any event ("orphan" markets) must be classified directly from their question text. To obtain a single category per market for the analyses in the paper, we build a hierarchical assignment procedure that uses event tags as the primary signal and successively falls back on weaker signals for events and markets that the earlier stages fail to classify. The goal is to place every market in exactly one of the seven categories (Sports, Crypto, Politics, Finance, Tech, Culture, Weather).

The procedure operates first at the event level. We start from the universe of Polymarket events, which contains 316,429 events covering 729,133 markets.

Stage 1: Event tags. We hand-curate a mapping from Polymarket tag identifiers to our seven categories. The map contains both top-level tags (for example, the "Politics" or "Sports" platform tags) and a larger set of more specific tags (for example, "NFL", "AI", or "Elections") that unambiguously belong to one of our categories. Tags that do not map to any of our seven categories are ignored. For each event, we take the tag list in the order supplied by Polymarket, identify every tag that has a category mapping, and assign the event to the category associated with the first such tag. When the same event has mapped tags from multiple categories, we also flag it as multi-category, so that later analyses can verify robustness to this choice. After Stage 1, 312,102 events (719,775 markets; 98.6% of all events) are categorized.

Stage 2: Legacy event category. For events whose tags contain no mapped entry, we fall back to the legacy "category" string on the event record and apply a second hand-curated mapping from those strings to our seven categories. This recovers events that were created before the tag taxonomy existed or whose tag list is empty. After Stage 2, 314,922 events (723,962 markets) are categorized.

Stage 3: Event-title heuristics. For events still uncategorized, we apply a set of hand-curated title rules: prefixes (e.g., titles starting with "Box Office:") and substrings (e.g., containing "$BTC" or "Will Trump") that reliably identify one of the seven categories. Matching is case-insensitive, and the first matching rule wins. After Stage 3, 316,180 events (728,001 markets) are categorized.

Stage 4: Manual event overrides. For the small residual set of events that none of the preceding stages classify, we assign categories manually by inspecting the event question. Stage 4 classifies the remaining 249 residual events, covering 502 markets, and brings the event-level total to 316,429 events (728,503 markets). The remaining 630 markets are orphan markets unlinked to any event, which we classify in the next step.

Market-level inheritance and orphan markets. Each market is then assigned the category of its parent event through the event identifier. However, 630 markets are not linked to any event in the Polymarket events dataset; we refer to these as orphan markets. For orphan markets, we run an analogous cascade directly on the market question text: first the event-title heuristics from Stage 3, then an additional set of market-question-specific rules, and finally a manual override list for the residual. After this step, every market in the sample is assigned to exactly one of the seven categories, yielding 729,133 categorized markets in total, and we drop no market on the basis of missing classification.

A.3. Determining the Winning Outcome

For every analysis that depends on whether a trader's position ultimately paid off, we need a reliable, token-level winner flag: each outcome token of each market should be labelled as a winner, a loser, or undetermined (for markets that never resolved cleanly). Polymarket exposes several fields that in principle encode the resolution outcome, but they are not all equally reliable. The raw market record includes an integer `outcome` field, an ordered list of outcome labels (e.g., "Yes"/"No" or two team names), an ordered list of final settlement prices (`outcomePrices`), and a resolution-status string (`umaResolutionStatus`) produced by the UMA oracle that adjudicates the market. Early validation of our pipeline revealed that the integer `outcome` field is not trustworthy, producing winner assignments that systematically disagree with the actual settlement outcome; we therefore do not use it. Instead, we rely on two complementary sources that together allow us to recover the winning token for the vast majority of resolved markets.

Primary source: simplified-markets winner flags. Polymarket publishes a "simplified markets" dataset that, for each market, lists the outcome tokens together with an explicit Boolean `winner` flag set by the resolution oracle. When a market is present in this dataset, we identify the winning token directly by the flag. We retain only markets for which exactly one token is flagged as the winner, discarding the small residual with zero or multiple flagged tokens as ambiguous.

Fallback source: settlement prices. A number of resolved markets are not yet present in the simplified-markets dataset. For those, we fall back on the `outcomePrices` field. Because this field is published for markets at all stages of their life cycle, including unresolved ones, we first restrict to markets whose `umaResolutionStatus` is either "resolved" or "settled", the two values that the oracle assigns after a market has gone through its validation window without an unresolved dispute. For those markets, settlement prices are normalized to one USDC (i.e., the winning token pays one USDC per share and the losing tokens pay zero). After rounding the reported prices to two decimals, we identify the outcome whose price equals one as the winner. Markets that do not contain exactly one such outcome (because all prices are strictly between zero and one, or because multiple outcomes round to one) are treated as ambiguous and left without a winner.

Combining the sources and labelling tokens. For each market, we take the simplified-markets signal when available and otherwise the settlement-price signal, yielding at most one winning outcome index per market. This single index is then propagated to the token level: the identified outcome's token is labelled a winner, every other outcome token of the same market is labelled a loser, and all tokens of markets for which no winner could be determined are labelled undetermined. The same procedure applies uniformly to binary and multi-outcome markets, including those issued under the negative-risk framework, because the one-USDC settlement rule holds regardless of the number of outcomes. The resulting token-level flag is the primary ingredient in our construction of realized payoffs and of the "bet on winner" indicator used throughout the analysis.

A.4. Computing Point-in-Time User Positions and PnL

A central input to the analysis is, for every wallet at well-defined points in time, a full snapshot of its token holdings, its USDC cash balance, and its implied marked-to-market PnL. Most of the analyses in the paper use each wallet's PnL at the end of the sample period, while a small number of figures rely on month-end snapshots. We construct this panel directly from the reconciled trade stream of Section A.1. The approach is conceptually simple: we translate each trade into a set of signed position changes for the two participating wallets, aggregate the changes to a daily frequency, forward-fill holdings between trading days, value open positions using prevailing market prices, and add the running cash balance to obtain PnL.

From trades to signed position changes. Each reconciled trade involves two wallets (a maker and a taker) and specifies the outcome token, the quantity in shares, the price in USDC per share, and the direction in which the taker traded. Because every trade settles against USDC, each trade also implies an offsetting cash flow. We therefore translate each trade into four position changes: two in the outcome token (one each for the maker and the taker, with opposite signs) and two in USDC (likewise with opposite signs), where the USDC amount equals price times quantity. We track USDC alongside outcome tokens, so that a wallet's complete state at any point in time is represented by a sparse vector of (token, quantity) pairs. No wallet is assumed to begin with an initial cash deposit: USDC balances evolve from zero purely as the net of the wallet's trading activity.

Daily aggregation and forward-fill. We aggregate position changes by (wallet, token, day) and cumulate them through time to obtain each wallet's end-of-day holding in each token. Between trading days, holdings are carried forward unchanged: a wallet that buys 100 shares on day t and does not trade again is treated as still holding 100 shares on day $t+k$ for all $k \geq 0$, until a subsequent trade or a market resolution changes the position. The same forward-fill logic applies to USDC.

Resolution and redemption. When a market resolves, its outcome tokens pay either one USDC per share (winner) or zero (loser), as described in Section A.3. We represent this as a single synthetic "settlement" event on the day following the market's close: for each wallet holding a token of the resolved market, the token position is closed out and the corresponding USDC amount (position times settlement price) is credited to the wallet's cash balance. Token positions are not forward-filled past the resolution date, since the token no longer trades.[22]

Marking to market. To value open positions on a given day, we use the last observed trade price for each outcome token at or before that day. Specifically, we construct a daily price series for each token from its last-traded price within each trading day and append a final row at the settlement price on the day of resolution. Open positions on a given snapshot day are then matched to prices through a backward point-in-time join, so a position that last traded several days earlier is valued at its most recent observed price, and a position in a resolved market is automatically valued at the final settlement price by the same mechanism.

The per-user panel and PnL. Each wallet's PnL on a given day is defined as the sum of (position $\times$ price) across all outcome tokens still open on that day, plus its running USDC balance. Because wallets start with zero cash and every trade's cash flow nets out within the trade, this quantity is directly interpretable as the wallet's realized-plus-unrealized profit (or loss) since its first trade. We construct this panel for every wallet, at daily frequency, from its first trade through the end of the sample, so that every active wallet has a snapshot for every day of its active life. To validate the construction, we stratify-sample wallets by activity level and compare our end-of-sample

[22]The blockchain does not technically block trades on outcome tokens after their market has resolved, and a small number of such trades do occur. We disregard these post-resolution trades in our computation of positions and PnL.

PnL to the values served by the Polymarket Data API; agreement is exact across all sampled wallets (see Section 3.2).

B. Wash Trading Robustness

The main text does not control for wash trading in any of its specifications. This appendix describes how we identify likely wash-trading wallets, reports summary statistics on their prevalence, and re-estimates the concentration and probit results after excluding them. The main findings of the paper are unaffected.

B.1. Identification procedure

For every wallet with at least 100 settled trades we compute the Herfindahl–Hirschman Index (HHI) of its counterparty distribution, where shares are the fraction of the user's total USDC volume transacted with each counterparty:

$$\mathrm{HHI}_i = \sum_c \left(\frac{\mathrm{vol}_{ic}}{\sum_{c'} \mathrm{vol}_{ic'}} \right)^2 .$$

A wallet that routes all of its volume through a single counterparty has $\mathrm{HHI}_i = 1$, while a wallet that splits volume uniformly across K counterparties has $\mathrm{HHI}_i = 1/K$. Self-dealing through a small ring of affiliated accounts mechanically produces a high HHI, so the measure is a natural proxy for the reciprocal-flow networks that characterize wash trading (Sirolly, Ma, Kanoria, and Sethi, 2025).

In what follows we flag a user as a suspected wash trader when $\mathrm{HHI}_i \geq 0.5$, i.e., when a single counterparty accounts for at least half of the user's USDC volume on its own (or a handful of counterparties jointly do so). The 100-trade screen removes casual users whose HHI is mechanically high because they trade only a handful of times. We report sensitivity to the 0.5 cutoff below.

This HHI-based rule is more conservative than the network-based detection procedure of Sirolly, Ma, Kanoria, and Sethi (2025), which can classify a trade as wash even when neither wallet is HHI-concentrated, provided the broader pattern of reciprocal flows fits a wash template. We therefore expect our flags to capture a smaller share of volume than the ∼25% they report. This cautious calibration is intentional: the goal of this appendix is to show that the paper's conclusions survive the removal of the users for whom the wash-trading interpretation is most defensible, rather than to estimate the total volume of artificial trading on the platform.

B.2. Summary statistics

Table B1 reports the distribution of counterparty HHI across the 482,092 wallets that meet the 100-trade screen. The distribution is highly skewed: the median wallet has an HHI of 0.032, the 95th percentile is 0.28, and the 99th percentile is 0.63. Applying the 0.5 cutoff flags 7,241 wallets (1.5% of screened users) jointly responsible for about $3.5 billion of their own volume.

Table B2 shows how the set of flagged wallets and the associated trade volume scale with the HHI cutoff. At the strict 0.9 threshold, 2,259 wallets (0.5% of screened users) are flagged, and trades involving at least one flagged wallet account for 2.0% of total platform volume. Relaxing the cutoff to 0.3 flags 20,804 wallets (4.3%) and captures 5.1% of volume. Even the loosest cutoff we consider is well below the ∼25% volume share reported by Sirolly, Ma, Kanoria, and Sethi (2025), consistent with our rule capturing only the most HHI-concentrated wallets rather than the full reciprocal-flow graph they analyze.

Table B3 reports the same decomposition by market category. Wash-adjacent volume is most prevalent in Sports (4.7% at the 0.5 cutoff) and Tech (4.0%), and smallest in Politics (2.3%) and Crypto (2.5%). Every category other than the small Untagged bucket lies well below the cross-category averages of Sirolly, Ma, Kanoria, and Sethi (2025).

Table B1: Counterparty HHI Distribution and Wash-Trading Flags

This table summarizes the distribution of counterparty Herfindahl–Hirschman Index (HHI) values across Polymarket wallets with at least 100 settled trades, and reports the number of wallets flagged as suspected wash traders under our baseline cutoff (HHI ≥ 0.5, volume-weighted). Three HHI variants are reported: Count (shares based on number of trades per counterparty), Volume (shares based on USDC volume), and Quantity (shares based on token quantity). Panel A reports the moments of each HHI variant. Panel B reports the number of flagged wallets and the volume of trades in which at least one side is a flagged wallet. The sample period is from November 11, 2022 to March 29, 2026.

Metric	Value
Panel A: HHI Distribution	
Users Analyzed	482,092
HHI (Count) — Mean	0.0284
HHI (Count) — Median	0.0151
HHI (Count) — P25	0.0098
HHI (Count) — P75	0.0253
HHI (Count) — P90	0.0478
HHI (Count) — P95	0.0837
HHI (Count) — P99	0.2637
HHI (Volume) — Mean	0.0726
HHI (Volume) — Median	0.0321
HHI (Volume) — P25	0.0172
HHI (Volume) — P75	0.0712
HHI (Volume) — P90	0.1816
HHI (Volume) — P95	0.2767
HHI (Volume) — P99	0.6284
HHI (Quantity) — Mean	0.0726
HHI (Quantity) — Median	0.0321
HHI (Quantity) — P25	0.0172
HHI (Quantity) — P75	0.0712
HHI (Quantity) — P90	0.1816
HHI (Quantity) — P95	0.2767
HHI (Quantity) — P99	0.6284
Panel B: Flagging Results	
Metric Used	HHI (Volume)
Threshold	0.5
Users Flagged	7,241
% of Users	1.5%
Volume of Flagged Users	$3,541,331,046
% of Volume	3.2%

Table B2: Wash-Trading Flags by HHI Threshold

This table reports the number and share of wallets flagged as suspected wash traders, and the number, share, and USDC volume of trades in which at least one side is a flagged wallet, at four volume-weighted counterparty HHI thresholds (0.3, 0.5, 0.7, 0.9). The screen and sample are those of Table B1.

HHI Threshold	Users Flagged	% Users	Trades	% Trades	Volume (USD)	% Volume
0.9	2,259	0.5%	471,564	0.1%	$1,374,183,019	2.0%
0.7	3,989	0.8%	826,938	0.1%	$1,661,426,121	2.5%
0.5	7,241	1.5%	1,488,118	0.3%	$2,083,649,435	3.1%
0.3	20,804	4.3%	4,339,481	0.7%	$3,397,918,594	5.1%

Table B3: Wash-Trading Volume Share by Market Category

This table reports, for each market category, the total USDC volume on the platform and the share of that volume in trades where at least one side is a wallet flagged under the counterparty-HHI rule, at four thresholds (0.3, 0.5, 0.7, 0.9). The screen and sample are those of Table B1.

Category	Total Volume	Wash % (HHI≥0.9)	Wash % (HHI≥0.7)	Wash % (HHI≥0.5)	Wash % (HHI≥0.3)
Sports	$25,477,038,956	3.4%	3.9%	4.6%	6.8%
Politics	$21,114,381,449	1.2%	1.6%	2.1%	4.1%
Crypto	$12,762,003,819	1.2%	1.6%	2.2%	3.3%
Culture	$3,473,852,981	1.0%	1.7%	2.8%	5.3%
Finance	$3,107,499,219	1.0%	1.4%	1.7%	4.7%
Tech	$806,902,466	2.6%	3.0%	3.3%	4.3%
Weather	$393,742,354	1.7%	2.0%	2.4%	3.5%
None	$5,320,659	0.0%	0.0%	0.0%	0.0%

B.3. PnL concentration excluding wash traders

Figure B1 replicates the main-text concentration snapshot (Figure 3) after dropping the 7,241 wallets flagged under the 0.5 cutoff. With wash traders removed, the 764,007 remaining winners share \$1.0 billion in aggregate gains; the top 0.1% capture 51.4%, the top 1% 76.7%, and the top 5% 91.5%—within a few tenths of a percentage point of the 51.2/76.5/91.4% reported in Section 4.1 for the full sample. Losses are likewise essentially unaffected: the bottom 0.1%, 1%, and 5% bear 43.9%, 68.2%, and 86.1% of total losses, versus 43.7%, 68.1%, and 86.1% in the main text. Figure B2 shows that the time-series pattern around the 2024 election and its subsequent stabilization are likewise unchanged. The extreme concentration documented in Section 4.1 is therefore not driven by wash-trading wallets.

Figure B1: Concentration of Gains and Losses, Wash Traders Excluded

This figure replicates Figure 3 after excluding the 7,241 wallets flagged as suspected wash traders under the baseline counterparty-HHI rule (HHI ≥ 0.5, volume-weighted; at least 100 settled trades). For each percentile threshold (Top 0.1%, 1%, 5%, 10%, 25%, 50%), the green bar reports the share of total gains captured by users in that top group (winners ranked from highest to lowest PnL), and the red bar reports the share of total losses borne by users in that bottom group (losers ranked from largest to smallest loss). The sample period is from November 11, 2022 to March 29, 2026.

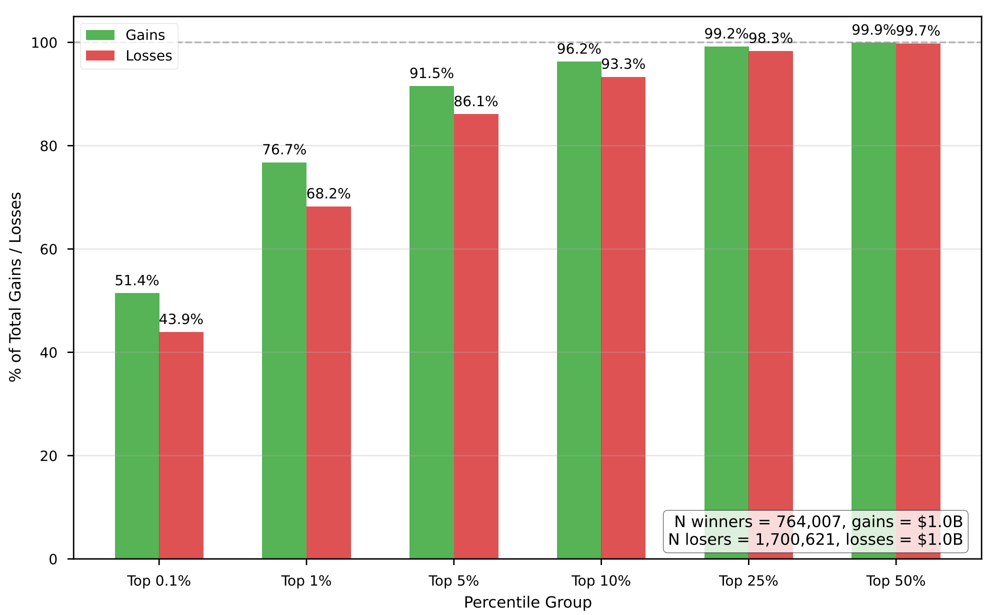

Figure B2: Concentration of Gains and Losses Over Time, Wash Traders Excluded

This figure replicates Figure 4 after excluding the 7,241 wallets flagged as suspected wash traders under the baseline counterparty-HHI rule. Panel A shows the concentration of gains and losses across Polymarket users over time for the top 0.1%, 1%, and 5% gains and losses. Panel B shows the fraction of users with positive PnL over time.

B.4. Loss-determinants probit excluding wash traders

Table B4 re-estimates the specifications of Table 5 on the wash-excluded sample. Entries are average marginal effects evaluated at the sample mean with robust standard errors, the same convention as the main-text table, so individual cells can be compared directly against their counterparts in Section 4.5. The overall story is that the marginal effects are essentially unchanged. Frac Extreme Price, Frac Maker Volume, Log N Trades, Log Total Volume, Category HHI, and each of the six category fixed effects added in column (4) match their main-text counterparts to the reported precision. In particular, the roughly 9.0 pp per 1 SD reduction in loss probability associated with Frac Maker Volume in column (4) and the sign reversal of Category HHI between column (3) and column (4) once category fixed effects are added both carry over unchanged.

The one coefficient that moves meaningfully is Counterparty HHI in the active-trader panels. Its main-text marginal effect of +18.0 pp for users with more than 100 trades (column (5)) becomes +35.4 pp once wash traders are removed, and the +50.0 pp main-text effect for users with more than 1,000 trades (column (7)) rises to +82.7 pp. Both remain highly significant. At first glance one might expect the opposite—the variable is marketed as a wash-trading proxy, and we have just dropped the wallets where it is most elevated. The direction of the change suggests that the mapping from Counterparty HHI to losing runs deeper than wash trading narrowly defined: in the residual active-trader sample, concentration in a small set of counterparties (e.g., repeated two-sided interaction with a dominant liquidity provider) is an even sharper marker of losing than in the full sample. The Frac Maker Volume and Frac Extreme Price findings, which carry the economic narrative of Section 4.5, are unaffected.

Table B4: Loss-Determinants Probit, Wash Traders Excluded

This table replicates Table 5 after excluding the 7,241 wallets flagged as suspected wash traders under the baseline counterparty-HHI rule. Entries are average marginal effects from the Probit specifications of Section 4.5, with robust standard errors in parentheses. The dependent variable is an indicator equal to one if the user's mark-to-market PnL is negative. The table reports the results for all users (columns 1–4), for users with more than 100 trades (columns 5–6), and for users with more than 1,000 trades (columns 7–8). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

	All users				Users with more than 100 trades		Users with more than 1,000 trades	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Frac Extreme Price	−0.012***	−0.017***	−0.007***	−0.003***	0.002	0.007***	−0.129***	−0.128***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.006)	(0.006)
Frac Maker Volume		−0.354***	−0.349***	−0.343***	−0.318***	−0.314***	−0.142***	−0.137***
		(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.006)	(0.006)
Log N Trades		−0.001***	0.004***	0.018***	0.030***	0.034***	−0.002	−0.005**
		(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.002)	(0.002)
Log Total Volume		0.020***	0.018***	0.020***	−0.022***	−0.023***	−0.044***	−0.041***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Category HHI			0.080***	−0.057***	−0.019***	−0.080***	−0.136***	−0.146***
			(0.001)	(0.002)	(0.003)	(0.004)	(0.007)	(0.010)
Counterparty HHI			−0.002	0.040***	0.354***	0.403***	0.827***	0.812***
			(0.002)	(0.002)	(0.010)	(0.011)	(0.063)	(0.063)
Traded Crypto				−0.074***		−0.018***		0.056***
				(0.001)		(0.002)		(0.006)
Traded Finance				−0.034***		−0.019***		−0.009*
				(0.001)		(0.002)		(0.005)
Traded Politics				−0.056***		−0.000		0.041***
				(0.001)		(0.002)		(0.005)
Traded Tech				−0.034***		−0.023***		−0.026***
				(0.001)		(0.002)		(0.005)
Traded Culture				0.002**		−0.005**		−0.023***
				(0.001)		(0.002)		(0.005)
Traded Weather				−0.036***		−0.024***		−0.008*
				(0.001)		(0.002)		(0.005)
Observations	2,472,860	2,472,860	2,472,860	2,472,860	470,964	470,964	78,223	78,223
Pseudo R^2	0.000	0.034	0.036	0.043	0.042	0.043	0.055	0.057

C. Calibration by Market Volume

Section 4 documents that Polymarket prices are well calibrated on average and breaks down calibration by market category. This appendix provides a complementary view by slicing the same sample on market trading volume. Panels are constructed as in Figure 5: prices are binned into ten equal-width bins, each panel's x-axis shows the bin midpoint and its y-axis shows the observed "Yes" frequency, the dashed 45-degree line marks perfect calibration, and shaded bands are 95% bootstrap confidence intervals (1,000 replications). Only markets with valid prices at all displayed horizons are included, yielding a consistent sample of 36,647 markets. Volume is measured as cumulative USDC traded before the 7-day reference horizon, and the same market-level bucket assignment is applied at every horizon.

Figure C3 splits markets into equal-count volume quintiles (roughly 7,329 markets each), so differences across panels cannot be attributed to sample size. The quintile cutoffs are $163, $1.0k, $11k, and $74k; panel Q1 therefore contains markets that had effectively no trade flow a week out, while panel Q5 is the top fifth of the sample by 7-day cumulative volume. In Q1 the observed frequency fluctuates well away from the diagonal and the bootstrap bands are wide, reflecting the fact that prices in a market with minimal trade flow carry little information. Calibration tightens progressively across Q2–Q4, and in Q5 both horizons hug the 45-degree reference closely with narrow bands. This is the pattern one would expect if trading volume, rather than market category per se, drives the cross-section of informativeness.

Figure C4 uses the fixed dollar cutoffs that match the Brier-score decomposition in Section 4, so panels line up directly with the corresponding Brier bars. The distribution of markets across buckets is top-heavy: the <$1k bucket alone contains 13,311 of the 36,647 markets in the consistent sample, while the $500k–1M bucket holds 969 and the $1M+ bucket 1,843. Confidence-band widths therefore reflect a combination of informativeness and sample size. Two features are worth highlighting. First, the low-volume panels (<\$1k through \$10k) have *tight* bands around clearly mis-calibrated point estimates: pooling thousands of illiquid markets reveals an S-shaped deviation from the diagonal whose qualitative shape echoes Q1 of the quintile figure. Second, the high-volume panels (>\$250k) have wide bands that straddle the diagonal: the underlying markets are individually well priced, but with at most a couple of thousand markets per bucket the bin-by-bin estimates are noisy. Reading the two figures jointly—equal-count quintiles for the cross-sectional pattern, fixed dollar buckets for comparability with the Brier decomposition—reinforces the same conclusion: Polymarket prices are best calibrated in the markets that attract real trading volume, and the overall calibration finding of Section 4 is driven by high-volume markets rather than by symmetric accuracy across the platform.

D. PnL Spread Decomposition by Category

The following tables replicate Table 5 separately for each market category. Each table reports mean PnL including and excluding the bid-ask spread by percentile bucket, for losers (Panel A) and winners (Panel B) within that category.

Figure C3: Calibration of Polymarket Prices by Volume Quintile

This figure plots the calibration of Polymarket binary market prices at different time horizons before resolution, with markets split into equal-count volume quintiles (Q1 lowest, Q5 highest) based on cumulative USDC volume at the 7-day reference horizon. The upper-left panel ("All") pools all quintiles for reference. Market prices are binned into ten equal-width bins (0–10¢, 10–20¢, . . . , 90¢–$1). For each bin, the x-axis shows the midpoint (expected probability) and the y-axis shows the observed "Yes" frequency. The dashed 45-degree line represents perfect calibration. Shaded bands show 95% bootstrap confidence intervals for the observed frequency (1,000 replications). Only markets with valid prices at all displayed horizons are included. Panel titles report the USDC range covered by each quintile. The sample period is from November 11, 2022 to March 29, 2026.

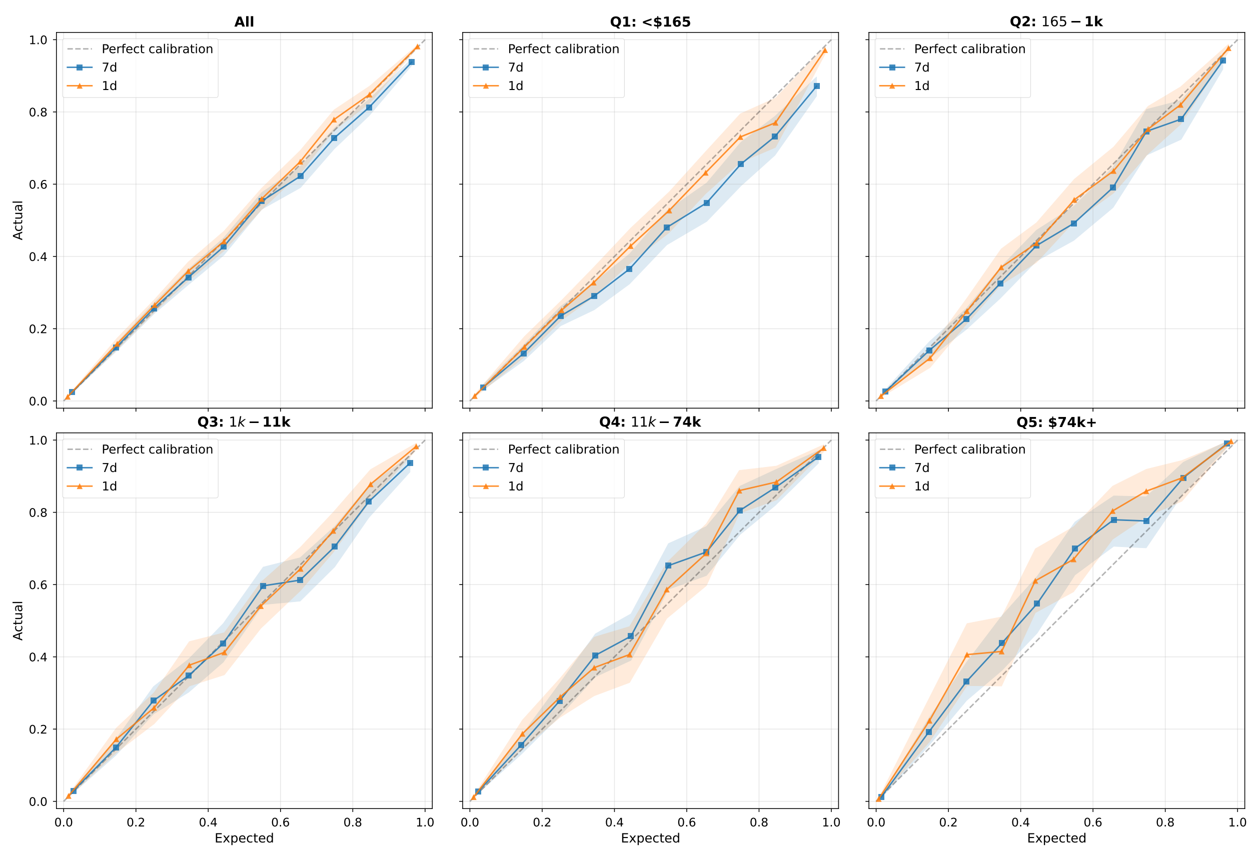

Figure C4: Calibration of Polymarket Prices by Volume Bucket

This figure plots the calibration of Polymarket binary market prices at different time horizons before resolution, with markets assigned to one of ten fixed volume buckets (<$1k, $1–5k, $5–10k, $10–25k, $25–50k, $50–100k, $100–250k, $250–500k, $500k–1M, $1M+) based on cumulative USDC volume at the 7-day reference horizon. The upper-left panel ("All") pools all buckets for reference. Market prices are binned into ten equal-width bins (0–10¢, 10–20¢, . . . , 90¢–$1). For each bin, the x-axis shows the midpoint (expected probability) and the y-axis shows the observed "Yes" frequency. The dashed 45-degree line represents perfect calibration. Shaded bands show 95% bootstrap confidence intervals for the observed frequency (1,000 replications). Only markets with valid prices at all displayed horizons are included; buckets containing fewer than ten markets are left blank. The bucket cutoffs match those used in the Brier-score decomposition of Section 4. The sample period is from November 11, 2022 to March 29, 2026.

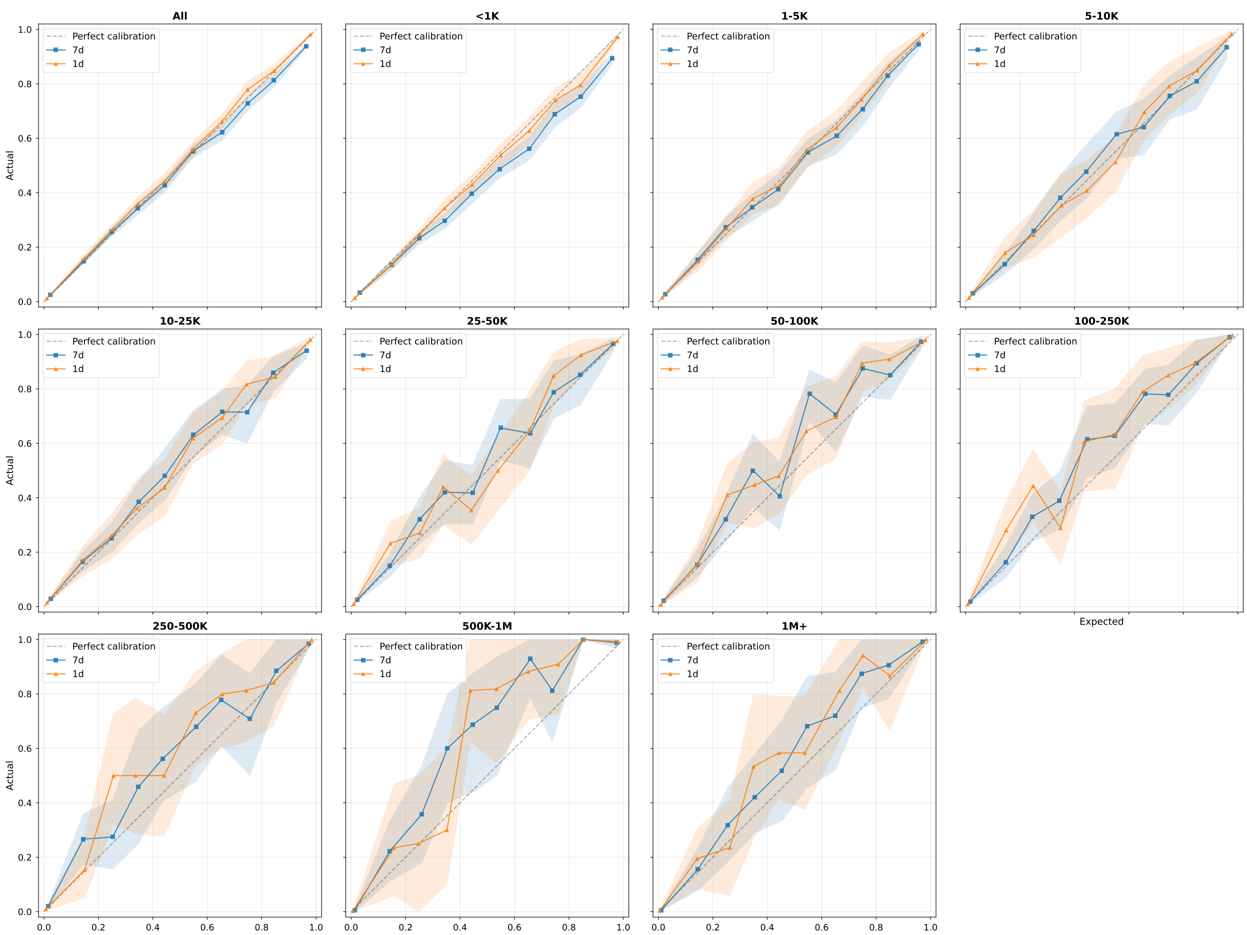

Table D5: PnL Spread Decomposition — Sports

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Sports** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL $<$ 0)						
0–20	211,417	-0.22	-0.00	-0.21^{***}	97.9%	15.4%
20–40	211,417	-1.88	-0.64	-1.24^{***}	65.9%	13.7%
40–60	211,417	-10.47	-7.88	-2.59^{***}	24.8%	15.5%
60–80	211,417	-55.70	-51.92	-3.78^{***}	6.8%	17.4%
80–90	105,709	-250.75	-243.25	-7.50^{***}	3.0%	17.0%
90–95	52,854	-862.53	-847.65	-14.87^{***}	1.7%	25.5%
95–99	42,284	-3,134.19	-3,090.97	-43.22^{***}	1.4%	28.7%
99–99.9	9,513	-20,143.64	-19,917.74	-225.90^{***}	1.1%	21.0%
99.9–100	1,057	-327,713.83	-323,742.90	$-3,970.93$	1.2%	27.5%
Panel B: Winners (PnL $>$ 0)						
0–20	99,122	0.08	0.20	-0.12^{***}	147.6%	23.5%
20–40	99,122	2.30	2.77	-0.47^{***}	20.3%	23.1%
40–60	99,122	13.04	14.16	-1.12^{***}	8.6%	24.2%
60–80	99,122	61.17	64.07	-2.90^{***}	4.7%	25.4%
80–90	49,561	293.38	298.27	-4.89^{***}	1.7%	28.7%
90–95	24,780	1,108.16	1,117.34	-9.17^{***}	0.8%	39.9%
95–99	19,824	4,045.72	4,035.42	10.30	-0.3%	47.9%
99–99.9	4,461	30,626.17	29,881.76	744.41^{***}	-2.4%	45.2%
99.9–100	495	595,602.95	577,803.98	$17,798.98^{***}$	-3.0%	54.3%

Table D6: PnL Spread Decomposition — Crypto

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Crypto** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL $<$ 0)						
0–20	154,359	-1.07	-0.65	-0.42^{***}	39.1%	17.7%
20–40	154,358	-10.19	-9.34	-0.85^{***}	8.3%	16.2%
40–60	154,358	-41.29	-39.63	-1.66^{***}	4.0%	16.6%
60–80	154,358	-183.41	-178.47	-4.93^{***}	2.7%	18.0%
80–90	77,179	-756.84	-741.82	-15.02^{***}	2.0%	19.2%
90–95	38,590	-2,029.71	-2,008.11	-21.61^{***}	1.1%	31.3%
95–99	30,872	-6,174.81	-6,116.71	-58.10^{***}	0.9%	21.8%
99–99.9	6,946	-27,862.56	-27,579.61	-282.95^{***}	1.0%	16.2%
99.9–100	771	-177,033.74	-175,888.43	$-1,145.31^{***}$	0.6%	15.4%
Panel B: Winners (PnL $>$ 0)						
0–20	116,813	0.10	0.14	-0.04^{***}	42.2%	23.4%
20–40	116,812	1.15	1.37	-0.22^{***}	18.7%	23.1%
40–60	116,813	7.08	7.42	-0.33^{***}	4.7%	24.6%
60–80	116,812	34.02	34.93	-0.91^{***}	2.7%	30.2%
80–90	58,406	165.75	167.36	-1.61^{***}	1.0%	37.8%
90–95	29,203	685.60	685.31	0.29	-0.0%	43.1%
95–99	23,363	4,146.07	4,097.22	48.85^{***}	-1.2%	50.1%
99–99.9	5,256	23,930.61	23,500.35	430.26^{***}	-1.8%	59.8%
99.9–100	584	217,583.35	209,525.89	$8,057.46^{***}$	-3.7%	58.4%

Table D7: PnL Spread Decomposition — Politics

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Politics** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL $<$ 0)						
0–20	196,042	-0.22	0.15	-0.37^{***}	169.2%	23.3%
20–40	196,041	-2.09	0.17	-2.27^{***}	108.2%	16.2%
40–60	196,042	-9.32	0.39	-9.70^{***}	104.1%	14.3%
60–80	196,041	-52.62	-43.33	-9.28^{***}	17.6%	13.5%
80–90	98,021	-275.15	-263.39	-11.76^{***}	4.3%	10.9%
90–95	49,010	-1,029.40	-1,012.26	-17.15^{***}	1.7%	11.8%
95–99	39,208	-4,885.85	-4,835.35	-50.50^{***}	1.0%	11.9%
99–99.9	8,822	-37,090.00	-37,008.49	-81.51	0.2%	11.9%
99.9–100	980	-441,610.79	-442,994.79	$1,384.00^{**}$	-0.3%	20.1%
Panel B: Winners (PnL $>$ 0)						
0–20	111,355	0.07	0.15	-0.08^{***}	121.6%	29.8%
20–40	111,354	0.88	1.23	-0.35^{***}	39.4%	23.9%
40–60	111,355	4.94	4.90	0.03	-0.7%	24.2%
60–80	111,354	26.94	25.30	1.64^{***}	-6.1%	26.3%
80–90	55,677	120.68	117.97	2.72^{***}	-2.3%	22.0%
90–95	27,839	444.49	438.66	5.84^{***}	-1.3%	23.0%
95–99	22,271	2,452.33	2,405.73	46.60^{***}	-1.9%	26.7%
99–99.9	5,011	23,752.33	23,407.33	345.00^{***}	-1.5%	35.1%
99.9–100	556	629,959.47	622,123.32	$7,836.15^{***}$	-1.2%	48.0%

Table D8: PnL Spread Decomposition — Culture

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Culture** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)						
0–20	83,820	-0.07	0.01	-0.09^{***}	118.9%	23.9%
20–40	83,819	-0.85	-0.34	-0.51^{***}	59.9%	17.8%
40–60	83,819	-5.26	-4.21	-1.05^{***}	19.9%	17.7%
60–80	83,819	-29.74	-28.57	-1.17^{***}	3.9%	20.3%
80–90	41,910	-132.13	-130.05	-2.08^{***}	1.6%	16.5%
90–95	20,955	-479.46	-475.83	-3.63^{**}	0.8%	15.1%
95–99	16,764	-2,349.76	-2,339.26	-10.50^{***}	0.4%	16.2%
99–99.9	3,771	-19,200.30	-19,188.44	-11.87	0.1%	11.6%
99.9–100	419	-263,776.71	-263,379.49	-397.22^{***}	0.2%	6.8%
Panel B: Winners (PnL > 0)						
0–20	56,615	0.03	0.07	-0.04^{***}	142.2%	27.2%
20–40	56,614	0.42	0.48	-0.06^{***}	14.5%	26.2%
40–60	56,614	2.97	3.06	-0.09^{***}	3.1%	25.4%
60–80	56,614	14.96	15.27	-0.31^{***}	2.1%	29.7%
80–90	28,307	54.76	55.24	-0.48^{***}	0.9%	29.6%
90–95	14,154	162.24	159.80	2.44^{***}	-1.5%	32.3%
95–99	11,323	880.81	859.72	21.09^{***}	-2.4%	36.4%
99–99.9	2,547	8,420.73	8,244.98	175.75^{***}	-2.1%	43.2%
99.9–100	283	164,904.64	164,522.62	382.02	-0.2%	61.7%

Table D9: PnL Spread Decomposition — Finance

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Finance** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL $<$ 0)						
0–20	72,454	-0.10	-0.03	-0.07^{***}	65.6%	21.5%
20–40	72,453	-1.06	-0.59	-0.47^{***}	44.1%	16.9%
40–60	72,454	-6.61	-5.74	-0.88^{***}	13.3%	17.5%
60–80	72,453	-33.02	-31.09	-1.92^{***}	5.8%	19.0%
80–90	36,227	-146.26	-143.62	-2.64^{***}	1.8%	18.0%
90–95	18,113	-529.71	-526.92	-2.79^{***}	0.5%	18.3%
95–99	14,491	-2,467.29	-2,456.23	-11.06^{***}	0.4%	20.3%
99–99.9	3,260	-18,861.03	-18,845.83	-15.20	0.1%	11.4%
99.9–100	362	-199,153.69	-199,183.68	29.99	-0.0%	11.5%
Panel B: Winners (PnL $>$ 0)						
0–20	67,086	0.02	0.07	-0.05^{**}	198.2%	30.7%
20–40	67,086	0.24	0.26	-0.02^{***}	9.2%	28.7%
40–60	67,085	1.89	1.98	-0.09^{***}	4.5%	23.5%
60–80	67,086	11.52	11.66	-0.15^{***}	1.3%	30.4%
80–90	33,543	50.87	51.11	-0.24	0.5%	31.2%
90–95	16,771	162.53	161.89	0.64	-0.4%	35.6%
95–99	13,417	881.57	871.08	10.49^{***}	-1.2%	39.0%
99–99.9	3,019	7,512.26	7,396.19	116.07^{***}	-1.5%	45.3%
99.9–100	335	171,619.83	171,276.06	343.77	-0.2%	61.9%

Table D10: PnL Spread Decomposition — Tech

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Tech** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)						
0–20	42,725	-0.08	-0.04	-0.04^{**}	52.3%	23.4%
20–40	42,725	-1.07	-0.87	-0.21^{***}	19.4%	18.9%
40–60	42,724	-6.57	-6.32	-0.25^{***}	3.8%	19.2%
60–80	42,725	-30.92	-30.40	-0.52^{***}	1.7%	22.1%
80–90	21,362	-118.29	-117.06	-1.24^{***}	1.0%	18.4%
90–95	10,681	-385.17	-381.45	-3.72^{***}	1.0%	17.9%
95–99	8,545	-1,663.14	-1,653.78	-9.37^{***}	0.6%	21.9%
99–99.9	1,923	-10,662.74	-10,648.21	-14.53	0.1%	17.5%
99.9–100	213	-98,723.51	-98,716.73	-6.79	0.0%	14.5%
Panel B: Winners (PnL > 0)						
0–20	47,256	0.02	0.02	0.00	-5.5%	26.4%
20–40	47,256	0.25	0.28	-0.04^{***}	14.7%	24.0%
40–60	47,255	1.52	1.55	-0.03^{**}	1.9%	22.6%
60–80	47,256	9.09	9.10	-0.01	0.1%	27.1%
80–90	23,628	33.47	33.61	-0.15^{*}	0.4%	31.7%
90–95	11,814	97.79	97.88	-0.09	0.1%	31.5%
95–99	9,451	498.98	492.38	6.61^{***}	-1.3%	35.3%
99–99.9	2,126	5,020.76	4,985.25	35.52^{**}	-0.7%	38.6%
99.9–100	236	56,911.03	56,532.54	378.49^{***}	-0.7%	58.3%

Table D11: PnL Spread Decomposition — Weather

This table reports mean profit and loss (PnL) including and excluding the bid-ask spread, by percentile bucket, for users who traded in the **Weather** category, separately for losers (Panel A: PnL < 0) and winners (Panel B: PnL > 0). See Table 5 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL $<$ 0)						
0–20	17,869	-0.39	-0.35	-0.04	9.8%	24.9%
20–40	17,869	-3.46	-3.16	-0.31^{***}	8.8%	21.4%
40–60	17,868	-13.81	-13.56	-0.25^{***}	1.8%	19.3%
60–80	17,869	-47.65	-46.87	-0.78^{***}	1.6%	19.2%
80–90	8,934	-152.55	-150.99	-1.57^{***}	1.0%	19.6%
90–95	4,467	-426.25	-422.41	-3.84^{***}	0.9%	18.6%
95–99	3,574	-1,612.57	-1,608.05	-4.52^{*}	0.3%	19.7%
99–99.9	804	-8,116.88	-8,097.49	-19.38^{**}	0.2%	17.3%
99.9–100	89	-52,959.27	-52,941.04	-18.23	0.0%	17.0%
Panel B: Winners (PnL $>$ 0)						
0–20	20,991	0.01	-0.01	0.02	-182.2%	26.6%
20–40	20,990	0.16	0.18	-0.02^{***}	9.6%	27.3%
40–60	20,990	1.29	1.36	-0.08^{***}	5.9%	22.5%
60–80	20,990	9.46	9.57	-0.11^{***}	1.1%	35.0%
80–90	10,495	34.02	34.26	-0.23^{***}	0.7%	33.2%
90–95	5,248	97.97	98.19	-0.22	0.2%	41.3%
95–99	4,198	581.23	575.09	6.13^{***}	-1.1%	49.2%
99–99.9	945	4,049.56	4,000.05	49.52^{***}	-1.2%	50.5%
99.9–100	104	50,547.73	50,308.76	238.96	-0.5%	58.1%

E. Resolved Markets Only

This appendix replicates three main-text tables on the subset of markets that had resolved by the sample end (March 29, 2026), i.e., markets for which a winning outcome is known on or before that date. Mark-to-market valuations of open positions on still-pending markets are excluded, so the resulting PnL reflects only realized payoffs from settled positions. Aside from the market-set restriction, the construction of every other variable follows the corresponding main-text table.

Table E12: PnL Concentration by Category — Resolved Markets Only

This table replicates Table 3 on the subset of markets that had resolved by March 29, 2026. Descriptive statistics and the distribution of mean PnL across percentile buckets are reported by market category. The top panel reports the number of markets, events, users, trades, and total volume for each category, as well as the fraction of users with positive PnL. The middle and bottom panels report mean PnL by percentile bucket for losers (PnL $<$ 0) and winners (PnL $>$ 0), respectively. Percentile buckets are constructed within each category independently. The sample period is from November 11, 2022 to March 29, 2026.

	All	Sports	Politics	Crypto	Culture	Finance	Tech	Weather
Descriptive Statistics								
N Markets	614,884	321,231	29,926	206,491	15,806	17,159	4,308	19,962
N Events	255,427	64,205	7,248	172,313	2,530	5,539	1,016	2,576
N Users (000s)	2,480.1	1,573.3	1,568.0	1,373.5	732.5	720.2	469.6	199.6
N Trades (M)	588.3	95.4	55.3	394.7	19.5	9.2	5.9	8.3
Volume (000s)	$67,140,742	$25,477,039	$21,114,381	$12,762,004	$3,473,853	$3,107,499	$806,902	$393,742
Frac. Winners	35.1%	33.1%	39.5%	42.2%	40.7%	47.7%	52.1%	52.5%
Mean PnL by Percentile – Losers (PnL $<$ 0)								
0–20	$-6.17	$-0.25	$-0.38	$-1.65	$-0.10	$-0.14	$-0.12	$-0.48
20–40	$-52	$-2.49	$-3.79	$-12	$-1.34	$-2.00	$-1.78	$-4.07
40–60	$-209	$-14	$-18	$-46	$-9.05	$-12	$-9.94	$-16
60–80	$-919	$-74	$-98	$-190	$-47	$-55	$-43	$-53
80–90	$-3,258	$-321	$-447	$-770	$-192	$-223	$-164	$-164
90–95	$-8,332	$-1,044	$-1,469	$-2,062	$-673	$-754	$-539	$-467
95–99	$-30,896	$-3,586	$-6,227	$-6,334	$-2,918	$-3,035	$-1,975	$-1,710
99–99.9	$-216,466	$-22,265	$-44,890	$-28,933	$-21,748	$-22,103	$-11,759	$-8,772
99.9–100	$-4,048,283	$-353,640	$-516,041	$-188,674	$-278,387	$-219,727	$-106,754	$-59,241
Mean PnL by Percentile – Winners (PnL $>$ 0)								
0–20	$1.92	$0.15	$0.16	$0.12	$0.04	$0.03	$0.02	$0.01
20–40	$23	$4.17	$2.33	$1.76	$0.68	$0.41	$0.31	$0.18
40–60	$108	$22	$14	$10	$5.12	$3.94	$2.31	$1.53
60–80	$549	$97	$71	$49	$25	$23	$13	$10
80–90	$2,282	$449	$289	$224	$94	$104	$51	$37
90–95	$6,251	$1,440	$919	$891	$294	$338	$155	$117
95–99	$24,564	$4,787	$3,745	$4,659	$1,416	$1,477	$787	$694
99–99.9	$216,811	$31,766	$27,710	$25,642	$10,090	$9,893	$5,606	$4,824
99.9–100	$4,854,299	$573,771	$637,746	$222,703	$177,627	$181,434	$59,259	$55,490

Table E13: Probit Regression: Determinants of Losses — Resolved Markets Only

This table replicates Table 5 on the subset of markets that had resolved by March 29, 2026. The dependent variable is an indicator equal to one if the user's PnL on the resolved-only subsample is negative. Reported values are marginal effects at the sample mean with robust standard errors. The independent variables and column structure follow Table 5. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

	All users				Users with more than 100 trades		Users with more than 1,000 trades	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Frac Extreme Price	−0.148***	−0.161***	−0.175***	−0.180***	−0.046***	−0.050***	0.047***	0.047***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.007)	(0.007)
Frac Maker Volume		−0.214***	−0.232***	−0.228***	−0.111***	−0.108***	0.075***	0.077***
		(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.006)	(0.006)
Log N Trades		0.010***	−0.014***	−0.018***	−0.028***	−0.023***	−0.010***	−0.010***
		(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.002)	(0.002)
Log Total Volume		0.021***	0.025***	0.023***	0.014***	0.009***	−0.001	−0.002
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Category HHI			−0.087***	0.002	−0.097***	−0.027***	−0.118***	−0.084***
			(0.001)	(0.002)	(0.003)	(0.004)	(0.008)	(0.011)
Counterparty HHI			−0.160***	−0.176***	0.020***	0.022***	0.263***	0.277***
			(0.002)	(0.002)	(0.007)	(0.007)	(0.042)	(0.042)
Traded Crypto				0.008***		−0.003		0.023***
				(0.001)		(0.002)		(0.007)
Traded Finance				0.016***		−0.006***		−0.013**
				(0.001)		(0.002)		(0.005)
Traded Politics				0.044***		0.060***		0.018***
				(0.001)		(0.002)		(0.005)
Traded Tech				0.019***		0.007***		0.021***
				(0.001)		(0.002)		(0.006)
Traded Culture				0.023***		0.016***		0.015***
				(0.001)		(0.002)		(0.005)
Traded Weather				0.008***		0.004**		−0.018***
				(0.001)		(0.002)		(0.005)
Observations	2,480,077	2,480,077	2,480,077	2,480,077	478,095	478,095	78,366	78,366
Pseudo R^2	0.010	0.026	0.032	0.034	0.007	0.009	0.006	0.007

F. No-Fee Markets Only

This appendix replicates the same three tables on the subset of markets that carry no platform taker fee (`taker_base_fee` null or zero in the Polymarket markets metadata). Markets in this subset account for roughly 85% of all markets in the sample; the excluded 15% are concentrated in daily Crypto "Up or Down" markets and NCAA Basketball matchups, which charge a 1% taker fee. PnL on the no-fee subset therefore isolates user performance in venues where transaction-cost differences across markets do not confound the cross-section.

Table F14: PnL Concentration by Category — No-Fee Markets Only

This table replicates Table 3 on the subset of markets without a platform taker fee. Descriptive statistics and the distribution of mean PnL across percentile buckets are reported by market category, with panels and bucket construction defined as in Table 3. The sample period is from November 11, 2022 to March 29, 2026.

	All	Sports	Politics	Crypto	Culture	Finance	Tech	Weather
Descriptive Statistics								
N Markets	614,884	321,231	29,926	206,491	15,806	17,159	4,308	19,962
N Events	255,427	64,205	7,248	172,313	2,530	5,539	1,016	2,576
N Users (000s)	2,480.1	1,573.3	1,568.0	1,373.5	732.5	720.2	469.6	199.6
N Trades (M)	588.3	95.4	55.3	394.7	19.5	9.2	5.9	8.3
Volume (000s)	$67,140,742	$25,477,039	$21,114,381	$12,762,004	$3,473,853	$3,107,499	$806,902	$393,742
Frac. Winners	34.2%	31.8%	35.5%	49.6%	38.6%	46.6%	50.3%	52.6%
Mean PnL by Percentile – Losers (PnL $<$ 0)								
0–20	$-6.66	$-0.22	$-0.22	$-0.81	$-0.07	$-0.10	$-0.08	$-0.39
20–40	$-61	$-1.86	$-2.09	$-8.33	$-0.85	$-1.07	$-1.08	$-3.47
40–60	$-244	$-10	$-9.31	$-34	$-5.26	$-6.63	$-6.57	$-14
60–80	$-1,046	$-55	$-53	$-161	$-30	$-33	$-31	$-48
80–90	$-3,594	$-249	$-275	$-748	$-132	$-147	$-118	$-153
90–95	$-9,207	$-858	$-1,029	$-2,161	$-480	$-531	$-386	$-427
95–99	$-33,795	$-3,135	$-4,887	$-7,131	$-2,351	$-2,472	$-1,665	$-1,613
99–99.9	$-224,950	$-20,391	$-37,095	$-44,966	$-19,200	$-18,888	$-10,682	$-8,117
99.9–100	$-3,131,290	$-365,960	$-441,858	$-1,018,961	$-263,777	$-199,154	$-98,724	$-52,959
Mean PnL by Percentile – Winners (PnL $>$ 0)								
0–20	$2.35	$0.08	$0.07	$0.14	$0.03	$0.02	$0.02	$0.01
20–40	$25	$2.42	$0.88	$2.19	$0.42	$0.24	$0.25	$0.16
40–60	$112	$14	$4.95	$14	$2.97	$1.88	$1.52	$1.30
60–80	$540	$64	$27	$95	$15	$12	$9.09	$9.50
80–90	$2,220	$310	$121	$599	$55	$51	$33	$34
90–95	$5,982	$1,158	$446	$2,362	$163	$163	$98	$98
95–99	$22,730	$4,242	$2,458	$9,856	$882	$888	$501	$581
99–99.9	$246,699	$32,954	$23,712	$55,846	$8,434	$7,537	$5,046	$4,051
99.9–100	$7,525,653	$633,831	$629,239	$499,640	$164,905	$171,593	$56,974	$50,539

Table F15: Probit Regression: Determinants of Losses — No-Fee Markets Only

This table replicates Table 5 on the subset of markets without a platform taker fee. The dependent variable is an indicator equal to one if the user's PnL on the no-fee subsample is negative. Reported values are marginal effects at the sample mean with robust standard errors. The independent variables and column structure follow Table 5. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

	All users				Users with more than 100 trades		Users with more than 1,000 trades	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Frac Extreme Price	−0.185***	−0.206***	−0.228***	−0.238***	−0.155***	−0.178***	−0.033***	−0.045***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.007)	(0.007)
Frac Maker Volume		−0.314***	−0.337***	−0.318***	−0.304***	−0.293***	−0.158***	−0.156***
		(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.006)	(0.006)
Log N Trades		0.012***	−0.014***	−0.006***	−0.029***	−0.006***	−0.029***	−0.013***
		(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.002)	(0.002)
Log Total Volume		0.029***	0.034***	0.029***	0.031***	0.013***	0.014***	−0.004***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Category HHI			−0.154***	−0.090***	−0.264***	−0.142***	−0.348***	−0.173***
			(0.001)	(0.002)	(0.003)	(0.004)	(0.008)	(0.011)
Counterparty HHI			−0.138***	−0.131***	−0.028***	0.017**	0.320***	0.356***
			(0.002)	(0.002)	(0.007)	(0.007)	(0.044)	(0.046)
Traded Crypto				−0.051***		−0.096***		−0.097***
				(0.001)		(0.002)		(0.007)
Traded Finance				−0.009***		0.010***		0.029***
				(0.001)		(0.002)		(0.005)
Traded Politics				0.077***		0.157***		0.138***
				(0.001)		(0.002)		(0.005)
Traded Tech				−0.012***		−0.018***		−0.034***
				(0.001)		(0.002)		(0.006)
Traded Culture				0.042***		0.026***		0.047***
				(0.001)		(0.002)		(0.005)
Traded Weather				−0.012***		0.002		0.013**
				(0.001)		(0.002)		(0.005)
Observations	2,480,077	2,480,077	2,480,077	2,480,077	478,095	478,095	78,366	78,366
Pseudo R^2	0.016	0.047	0.055	0.061	0.041	0.056	0.035	0.049