

Who Wins and Who Loses In Prediction Markets? Evidence from Polymarket*

Pat Akey¹, Vincent Grégoire², Nicolas Harvie³, Charles Martineau³

¹ESSEC Business School, CEPR, ECGI

²HEC Montréal

³University of Toronto

May 24, 2026

We study trading gains and losses on Polymarket, the largest prediction market. Using 588 million trades (\$67 billion in volume), we show that the gains are highly concentrated: the top 1% of users capture 76.5% of profits. Successful traders provide liquidity using limit orders that resolve favorably relative to realized outcomes while unsuccessful traders take liquidity using market orders. Monthly performance is weakly persistent, however, this may represent sample selection rather than skill. A detailed analysis of the trading behavior of the most successful accounts suggests that “insider” trading is unlikely to explain the performance of the largest winners.

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 Católica de Chile, Aarhus University, Wolfe Research 3rd Annual Canadian Quantitative and Macro Investment Conference, and 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, Rotman FinHub, and the Ontario Securities Commission. The data for this study is available at <https://huggingface.co/datasets/vgregoire/polymarket-users>. 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, with platforms like Polymarket and Kalshi allowing users to trade on the outcomes of real-world events spanning politics, economics, sports, and entertainment. These platforms are now routinely featured in major news outlets and have struck partnerships with media organizations that cite their contract prices alongside polls and analyst forecasts when covering political, economic, and sporting events.¹ While prediction markets are not new, innovations in fintech and blockchain technologies have allowed them to scale rapidly, with trading volume reaching \$67 billion USD on Polymarket and \$53 billion on Kalshi² in the last four years. Their rapid growth has prompted media commentators and policy makers to raise questions about the gamification of trading and the role of prediction markets as an entry point into financial markets, and concerns have also been raised about insider trading on these platforms.³

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 wealth accumulation (Odean, 1999; 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?

¹See, for example, <https://www.businessinsider.com/kalshi-cnbc-deal-cnn-data-integration-partnership-2025-12>

²Statistic for Kalshi are retrieved from <https://www.kalshidata.com/>. Polymarket statistics are retrieved from our analysis.

³See, 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/>.

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, 69.0% of users lose money. Second, Polymarket prices are well-calibrated on aggregate. A contract priced at p resolves in its favor approximately p percent of the time but this aggregate is concentrated in high-volume markets and high-attention categories: low-volume markets show deviations from perfect calibration, and the Tech, Culture, and Weather categories deviate even at the 1-day horizon. Third, the excess hit rate, defined as the average of realized outcomes minus trade prices, declines sharply through the PnL distribution, from roughly 20 percentage points in the most profitable percentiles to between -15 and -18 percentage points across the loss tail. 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. Moreover, we find that month-to-month performance is only weakly persistent across traders and show that there is a large amount of turnover in which traders continue to participate in Polymarket. Fourth, 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 \(2009, 2014\)](#) in Taiwanese equity markets and suggests that investor sophistication is an important predictor of the cross-sectional distribution of winners and losers on Polymarket. Fifth, we examine the trading behavior of the 100 most successful users in our sample (representing 27.7 percent of aggregate winners’ profits on

Polymarket) to understand whether the gains are generally consistent with informed trading and, in particular, with insider trading.⁴ We find that the most successful users traded frequently in sports markets, often for different teams (81% of the gains), and markets related to the 2024 US election (representing 15% of the gains). Within this group, gains are typically concentrated in a small number of markets—most top earners derive the bulk of their profit from one or a handful of markets rather than from a broad portfolio—and a meaningful fraction earn the bulk of their profit not from directional bets at all but from liquidity provision: posting many limit orders and capturing the spread on each at high volume. Given the nature of these trades, it seems unlikely that the most successful traders were benefiting from insider information, as opposed to forecasting skill—whether through quantitative modeling or deep event-specific expertise—or skilled liquidity provision. For insider information to be the source of profits for these traders it would need to be the case that the same trader had such information for many different sports teams spanning many sports or that they had some source of information that determined the outcome of millions of American voters before they had actually begun casting votes for the 2024 US presidential election. While we do not claim that there is *no* insider trading, the 100 most successful traders do not seem to have earned large profits primarily because of such trading.

Our results are robust to various alternative specifications and subsample analyses. We show that our main findings are not driven by wash trading, which is a concern in this setting given the ease of creating multiple accounts and trading between them. We further replicate our concentration results on the subset of fully resolved markets, where PnL reflects realized payoffs only rather than mark-to-market valuations, and on the subset of no-fee markets, where differences in platform transaction costs cannot confound the cross-section of performance. Both subsamples yield similar

⁴Throughout the paper, we use *insider trading* to refer to trading on non-public material information about a resolving event, as distinct from *informed trading*, which we use more broadly to describe any trade in which a participant holds superior information about an outcome, whether from a quantitative forecasting model, domain or event-specific expertise, or other legitimate sources.

findings. We also show that our persistence results are robust to risk adjustment. Specifically, we conduct performance analysis using a modified measure of the Sharpe and Sortino ratios and find similar evidence of weak month-to-month persistence. A further caveat is that 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.

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. In order to facilitate further study of these novel and increasingly important markets, we have made the full data used in this paper publicly available for all researchers.⁵

To clarify how aggregate price efficiency and extreme profit concentration can arise jointly, we complement the empirical analysis with a simulation benchmark. The model is a parsimonious repeated binary-market economy in which traders differ in initial wealth, true and perceived information precision, and propensity to provide liquidity. Prices are generated from a public log-odds signal that is unbiased on average but contains small residual mispricings. We note that we do not try to “match” moments of the empirical distribution, but rather attempt to provide a structured “null” hypothesis to better allow readers to interpret our findings. Each of the three behavioral and informational ingredients in the model—heterogeneous private signal precision, overconfidence,

⁵The data is available at <https://huggingface.co/datasets/vgregoire/polymarket-users>.

and heterogeneous maker propensity—contributes to the simulated cross-sectional dispersion of dollar PnL. Shutting any one of them down reduces simulated top-share concentration relative to the baseline. Turning all three on jointly, with the four trader primitives drawn independently, moves the simulated top-share concentration of gains and the excess-hit-rate gradient toward their empirical counterparts.

For policymakers, our research 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., Kumar, 2009; Dorn, Dorn, and Sengmueller, 2015; 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).⁶ 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

⁶Yeola, Allen, Desai, Poliak, Yang, Smith, and Ayers (2025) document a 23% rise in gambling-addiction help-seeking following sports-betting legalization.

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 paper to ours 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. Similar heterogeneity in trader profits has been documented in options markets ([Hu, Kirilova, Park, and Ryu, 2024](#)) and in commodity futures markets ([Dewally, Ederington, and Fernando, 2013](#)).

A second strand of literature documents behavioral anomalies and efficiency failures in prediction markets. [Snowberg, Wolfers, and Zitzewitz \(2013\)](#) 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 \(2026\)](#) 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. [Ng, Peng, Tao, and Zhou \(2026\)](#) compare Polymarket, Kalshi, PredictIt, and Robinhood during the 2024 election and find that more liquid prediction markets outperform polls, that Polymarket leads Kalshi in price discovery, and that net order imbalance from large trades predicts subsequent returns. [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.⁷ Another strand of research finds that prediction market prices provide information comparable to that of professional forecasters ([Chernov, Elenev, and Song, 2026](#); [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

⁷[Cookson, Niessner, and Schiller \(2026\)](#) provide evidence that social media can inform corporate decisions; an open question is whether prediction markets can play a similar role for corporate decision and investor behavior.

platform with full transaction transparency. [Reichenbach and Walther \(2025\)](#) study skill persistence and forecasting accuracy. [Della Vedova \(2026\)](#) classifies users as human or bot using an ad-hoc approach, decomposes profitability into directional and execution components, and finds that execution timing rather than forecasting ability is the dominant driver of dollar profits. 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. We also relate a richer set of user-level behavioral characteristics, including activity scale, longshot exposure, and category breadth, to the probability of loss. To our knowledge, no study documents the full distribution of profits and losses across participant types in any prediction market.

Finally, [Gomez Cram, Guo, Jensen, and Kung \(2026\)](#) similarly find that profits on Polymarket are concentrated and conclude that these traders are informed. We argue that evaluating such claims is complicated without accounting for the fact that there is very likely selection in which traders continue to trade and those that do not. Our results suggest that 44 percent of traders stop trading after a month (up to 66 percent after six months including 55 percent of the best performing traders). We find little evidence that successful (unsuccessful) traders exhibit substantial persistence across time (conditional on continuing to trade). We note that tests of dollar profits are difficult to adjust for risk given the short-term nature of the contracts.⁸ Additionally, our results highlight that successful liquidity provision plays an important role in the performance of the highest earning users, and find little evidence that the top 100 users traded in ways that we believe to be consistent with insider trading. We encourage readers to interpret claims about trader skill, informed trading, and the role of market making in both studies with these limitations in mind.

⁸This critique applies to our analysis as well, although we do estimate our persistence tests using Sharpe- and Sortino-analog ratios adapted to dollar profit measures.

2. Institutional Background

Polymarket is a decentralized prediction market platform, while Kalshi is a CFTC-regulated centralized event-contract exchange. Both let 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.⁹ 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.¹⁰ 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

⁹The actual smart contract structure allows for markets with more than two outcome 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.

¹⁰Not 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.

(CLOB) matching engine with on-chain settlement of conditional tokens.¹¹ 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 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 \$0.001.¹² However, certain markets have taker fees to fund the Maker Rebates Program and to incentivize deeper liquidity and tighter spreads.¹³ We remove these markets from our analysis in robustness checks and find that our main results are unchanged.

¹¹The 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.

¹²Note 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.

¹³The 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).

Contract resolution on Polymarket is determined by an oracle, a decentralized mechanism that reports real-world outcomes to the blockchain.¹⁴ 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.¹⁵

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.¹⁶ 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

¹⁴Polymarket uses the Universal Market Access (UMA) Protocol as its resolution oracle.

¹⁵For 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.

¹⁶Polymarket’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 the Polygon blockchain is necessary to obtain a complete sample.

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 Section A of the Internet Appendix.

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., 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. For Yes/No markets we 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 $1 - p$. For other two-outcome markets (such as head-to-head matchups), we normalize to the first outcome by index, which lets us include these markets in the accuracy analysis without imposing a Yes/No restriction. We describe the procedure that we use to identify the winning outcome for each market in Section A.2 of the Internet Appendix.

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. We also observe only Polymarket activity, so for sophisticated users who trade simultaneously on other prediction-market platforms such as Kalshi or PredictIt, our PnL captures only the Polymarket leg of any cross-platform arbitrage strategies they may pursue.

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 $\text{PnL} = \text{portfolio value} + \text{net USDC balance}$, meaning the USDC balance tracks realized gains and losses while the portfolio value tracks unrealized ones.¹⁷ 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.¹⁸ Positions are forward-filled from a user’s first trade through March 29, 2026, so every 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 Section A.3 of the Internet Appendix.

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

¹⁷Our 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.

¹⁸The net USDC balance we report is therefore a trading-only quantity: external on-chain USDC deposits and withdrawals are not netted into the construction. The PnL we compute equals the user’s portfolio profit from trading on Polymarket and is independent of any external cash flows.

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, while the mean is substantially higher (474 trades, 54 markets), reflecting a long right tail of highly active participants. Notional volume is similarly dispersed: the median user trades \$2,343, the mean is \$54,144, and the 95th percentile reaches \$97,793. A large share of trading occurs at extreme prices below 10¢ or above 90¢: the mean fraction is 54.9% and the median is 58.3%, with substantial cross-sectional variation (standard deviation of 39.7%). 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 and the mean is \$0, 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.

Panel (c) of Figure 1 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.¹⁹ 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.

3.4. 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. Figure 2 addresses this directly by plotting the realized frequency of winning outcomes against price bins using prices 7 days and 1 day before resolution.²⁰ 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. The first panel of Figure 2 shows that across the full price range and across all markets, 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.

¹⁹This 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. Section A.4 of the Internet Appendix details the full assignment procedure.

²⁰Section B of the Internet Appendix reports the calibration results for different time horizons up to 90 days before resolution.

The remaining panels of Figure 2 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 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 closer to 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.

4. The Cross-Section of Profits

In this section, we examine how trading gains and losses are distributed across Polymarket users and how that distribution lines up with a per-trade measure of forecasting accuracy.

4.1. Concentration of gains and losses

We next examine the concentration of gains and losses across Polymarket users. Focusing on dollar profits and losses, rather than on returns, allows us to more accurately characterize the distribution of winners and losers. For example, a contract trading at \$0.01 that resolves to zero gives a -100% return; one trading at \$0.99 that resolves to 100% gives a return of approximately 1%. Hence, returns cannot be aggregated in a meaningful way to measure profitability. Additionally, [Cziraki and Gider \(2021\)](#) show that the correlation between dollar profits and percentage returns is modest because returns are negatively correlated with trade size.

Panel (a) of 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. Figure 3 also shows that losses are 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%.

Panel (b) of Figure 3 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. Table 3 shows that the concentration of gains and losses is similarly extreme across all market categories, with the largest dollar gains for the top 0.1% concentrated in the Sports, Politics, and Crypto categories.

4.2. 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 value-weighted excess hit rate (EHR_{vw}), which we compute for each user i as:

$$\text{EHR}_{\text{vw},i} = \sum_k q_k(o_k - p_k) / \sum_k q_k, \quad (1)$$

where q_k is the traded quantity in shares for trade 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.²¹ 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.²²

Figure 4 plots the volume-weighted excess hit rate EHR_{vw} across the full PnL distribution. We rank users by PnL in ascending order and assign them to 200 half-percentile bins, with percentile 0 (least profitable) on the left and percentile 100 (most profitable) on the right. For each bin we plot the mean (solid line, with 95% confidence band) and the median (dashed line) of EHR_{vw} . The horizontal dashed line marks the sample mean across all users. Vertical dotted lines mark the bins containing the PnL thresholds of $-\$100$, $\$0$, and $\$100$, which fall near percentiles 15, 70, and 92, 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 $\$100$.

The figure shows that mean EHR_{vw} is strongly negative across the loss tail, falling to roughly -0.15 to -0.18 for percentiles below the $\$0$ threshold (percentile 70), before increasing to zero in the breakeven region and increasing more sharply near the right edge of the distribution, reaching about $+0.20$ for the most profitable users. 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 (dashed line) tracks the mean’s sign but is smaller in magnitude across

²¹Buy-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.

²²Under 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 EHR_{vw} .

the distribution, reflecting the long-tailed cross-section of per-user EHR_{vw} : a minority of traders in each bin drive the extremes of the mean. The fact that the volume-weighted metric is large in absolute value at both ends indicates that both winners and losers tilt their largest trades in a more extreme 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.

5. Who Wins and Who Loses?

To identify determinants of negative trading outcomes, Table 4 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.²³ Column (4) adds 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. 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 4 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).

Frac Maker Volume has the largest economic effect. 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 compared to other behavioral predictors. The magnitude 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. As in our previous performance tests, we caution that this cross-sectional association does not establish a causal role for market-making.

²³The 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 C in the Internet Appendix describes our HHI-based identification procedure, reports summary statistics on flagged wallets and their volume, and replicates our main analysis with flagged wallets excluded. None of the paper’s main findings are materially affected.

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 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 losers trade at extreme prices roughly twice as often as the top earners (Table 5).

Log N Trades has a small negative marginal effect in column (2) (-0.1 pp) but turns positive once we control for categories, 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 excess trading reflecting 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.²⁴ 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

²⁴The two variables are moderately correlated (0.65) and together control for the scale of trading activity.

broader population.

The *Category HHI* coefficient is informative about users' market category specialization. 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 volume to a single topic without any 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 users focusing on markets they know well. We find the same effect among users with more trading, where higher concentration predicts lower loss probabilities throughout (columns 5–8) and the addition of category fixed effects in column (6) increases the magnitude from -1.0 to -6.7 pp.

The *Counterparty HHI* estimate suggests that trading with the same set of users increases the likelihood of losses. 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 increases among more active traders, approximately $+20.6$ pp for users with more than 100 trades (column 6) and $+49.4$ pp for those with more than 1,000 trades (column 8).

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 close to zero ($+0.2$ 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.²⁵ Once category fixed effects are included, within-category specialization is associated with lower loss probabilities, and users whose trade volume is predominantly in Sports are not distinctive performers in the full sample but fare relatively better among the most active traders. We further report in the Internet Appendix, Sections D and E, that these results are robust to restricting the sample to only resolved markets²⁶ or only markets with zero transaction fees, respectively.

Panel A of Table 5 provides a descriptive complement to the probit results by directly comparing the trading characteristics of the most profitable users against the bottom 95%. The top 0.1% of earners supply 47.2% of their volume as makers, 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. 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. The activity gap is the largest dimension of all: the top 0.1% execute on average 150,651 trades each, against just 222 for the bottom 95%, consistent with a small group of high-frequency professional traders coexisting

²⁵Counterparty identities on Polymarket are only observable ex post, after a trade settles on-chain. The Polymarket order book displays the total quantity available at each price level but does not reveal individual resting limit orders or the wallet IDs behind them, so concentration on a small set of counterparties is not the result of users deliberately routing their orders to a known wallet.

²⁶The only material difference is for the group of the most active traders, which have a positive sign for both *Frac Extreme Price* and *Frac Maker Volume*

with a much larger population of casual users. Together, our results suggest that winning traders are sophisticated liquidity providers, while losing traders are liquidity takers who repeatedly overpay for low-probability outcomes.

Panel B of Table 5 reports risk-adjusted performance metrics computed at the user level from each user’s monthly mark-to-market PnL series, following the monthly aggregation convention used in Section 6.1. We construct two reward-to-risk ratios as analogs to the Sharpe and Sortino ratios but computed on dollar PnL rather than percentage returns: the *annualized PnL-to-volatility ratio*, defined as $\sqrt{12} \times \overline{\Delta \text{PnL}} / \sigma(\Delta \text{PnL})$, and the *annualized PnL-to-downside-volatility ratio*, defined analogously with the denominator replaced by the root-mean-square of negative monthly PnL changes. We avoid the labels “Sharpe” and “Sortino” because the standard definitions require returns, and constructing returns on Polymarket is not straightforward: traders can deposit and withdraw freely and there is no natural notional capital base.²⁷ The top 0.1% of earners attain a median annualized PnL-to-volatility ratio of 1.73, the next 0.9% reach 1.47, and the next 4% reach 1.46, while the bottom 95% have a negative median ratio of -1.15 . The PnL-to-downside-volatility ratios are similarly higher for all three top groups than for the bottom 95%. A PnL-to-volatility ratio near 1.7 for the top 0.1% is, on a Sharpe-like scale, in the range that characterizes skilled discretionary trading rather than the level one would expect from concentrated, uncompensated bets on long-shot binary outcomes, and the top group’s mean monthly PnL of \$54,607 compares to a maximum dollar drawdown of \$35,212—less than one month of expected profit—further indicating that their cumulative gains compound smoothly across their tenure rather than arising from a small number of realized tail outcomes. These statistics provide complementary support for the skill

²⁷A user-month is included if the user either traded in the month or carried a non-zero, non-USDC position into it from the prior month: months with no platform exposure have no economically meaningful PnL change. The two ratios are reported as cross-sectional medians within each group because per-user ratios have a heavy upper tail, i.e., a few highly profitable users with small downside semi-deviations produce extreme PnL-to-downside-volatility values, so the cross-sectional mean is dominated by those users while the median describes the typical user. Per-user ratios are computed only for users with at least three monthly observations and the relevant standard deviation above one dollar; the level rows (mean PnL, SD, max drawdown) are user-mean averages within each group.

interpretation that does not rely on month-to-month persistence: even though individual monthly outcomes are noisy in absolute dollars (top-group monthly SD of \$104,098), the longer-horizon PnL-to-volatility ratio is large enough to be informative about skill in a way that the monthly transitions of Section 6.1 cannot be. The bottom row of Panel B reports the cross-sectional mean of each user’s excess hit rate, which captures forecasting accuracy at the per-trade level: top traders exhibit substantially higher excess hit rates (0.097 for the top 0.1% against -0.027 for the bottom 95%), reinforcing that the risk-adjusted profitability gap above is paired with a measurable difference in the ability to identify mispriced contracts.

6. Forecasting Skill and Liquidity Provision

6.1. 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 into eleven groups based on the PnL they realized during a given month,²⁸ using the dollar breakpoints $-100k$, $-10k$, $-1k$, -100 , -10 , 10 , 100 , $1k$, $10k$, and $100k$. We then examine which group each user falls into 1 and 3 months later.²⁹ The bucket-to-bucket cells report transition frequencies *conditional on trading* (raw probability divided by one minus the No Trade rate); the “ > 0 ” column reports the conditional probability that the trader’s target-month PnL change is strictly positive, with significance from a two-sided exact-binomial test of the null hypothesis that the probability of profit equals 0.5; and the “No

²⁸Monthly 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.

²⁹Note 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 3-month performance looks at the PnL in April, not the cumulative PnL from February to April. Internet Appendix Table F10 reports the corresponding 6-month and 12-month transitions.

Trade” column reports the unconditional share of users with no positions in the target month.

Panel A of Table 6 presents results of this analysis at the one-month horizon. Looking first at the “No Trade” column—which records the share of users in each starting bucket who took no positions in the target month, neither trading nor carrying an open position into it—the five loss buckets show one-month no-trade rates clustered between 34% and 45% (36.8% for $<-100k$, 34.3% for $[-100k, -10k]$, 41.9% for $[-10k, -1k]$, 45.4% for $[-1k, -100]$, and 42.7% for $[-100, -10]$), while the best performers (the $>100k$ group) sit out only 16.7% of the time, with non-participation declining from the small-gain buckets toward the extreme winners. We do not observe deposits, withdrawals, or account closures, so we cannot tell whether these users have left the platform or simply paused; what we can say is that lower performance positively predicts subsequent non-participation, though even the most profitable traders take no positions in a non-trivial share of the following months. Conditional on continuing to trade, the two extreme tails are highly volatile rather than persistent: in the worst-loss group ($<-100k$), only 18.1% remain in the same bucket while 21.5% jump to $(10k, 100k]$ and 16.8% jump to $>100k$; in the highest-gain group ($>100k$), only 29.8% stay in the same bucket while 17.8% fall to $<-100k$ and 15.4% to $[-100k, -10k]$. The “ > 0 ” column corroborates this asymmetry while pinning down where the tails are noisy: the four interior loss buckets ($[-100k, -10k]$ through $[-100, -10]$) record 39–44% chances of a positive PnL change next month, all significantly below 50% at the 1% level, but the most extreme losers ($<-100k$) are at 48.7%—statistically indistinguishable from a coin flip ($n = 560$ traders). On the gain side, all gain buckets except $(100, 1k]$ (which lands exactly at 50.0%) lie significantly above 50%, with the two highest gain buckets reaching 60.5% and 59.0%. The “breakeven” traders (those with PnL change in $[-10, 10]$) exhibit a strong likelihood of remaining in the same group: conditional on trading, they have an 89.3% likelihood of having the same level of profits one period later, and they record a 57.4% probability of a positive PnL change. We observe qualitatively similar patterns at the three-month

horizon, where every starting bucket including $< -100k$ now reads significantly different from 50% as the small-sample noise washes out. Thus, while we find some evidence that a small group of traders have higher than “expected” hit rates, monthly performance does not persist through time, which challenges the interpretation that many traders have substantial skill.

Table 7 re-runs this exercise after replacing the dollar-PnL bucketing of Table 6 with an annualized risk-adjusted ratio — a Sharpe analog computed at the daily frequency. For each user-month with at least 10 active days, we compute the mean and standard deviation of daily PnL changes, take their ratio, annualize by $\sqrt{365}$ (since Polymarket trades seven days a week), and bucket user-months into nine groups with cutoffs at ± 10 , ± 5 , ± 3 , and ± 1 , leaving a single merged middle bucket $[-1, 1]$ that absorbs the small-magnitude ratios near breakeven.³⁰ Conditional on continuing to trade, the bucket-to-bucket transitions show the same qualitative pattern as Table 6: the two extreme tails are highly volatile rather than persistent, and the conditional probability of a strictly positive next-month ratio is well below 50% across the loss buckets and well above 50% across the gain buckets but does not jump to values that would indicate strong individual-level persistence. Risk-adjusting away the dollar-magnitude scale, in other words, does not reveal hidden month-to-month persistence: the weak-persistence finding of Table 6 is not an artifact of the dollar-PnL bucket boundaries. The Sortino-analog version (which replaces the volatility denominator with the downside semi-deviation) and the corresponding 6- and 12-month long-horizon variants are reported in Internet Appendix Section G and tell the same story.

³⁰User-months whose daily-PnL standard deviation is at or below \$0.10 are dropped, since the resulting ratio is dominated by floating-point noise on near-flat panels. As with the user-level Sharpe and Sortino analogs reported in Table 5, the ratio is computed on dollar PnL rather than percentage returns and is therefore not a strict Sharpe ratio.

6.2. Spread decomposition and liquidity provision

We now examine the role of liquidity provision and spreads on trader profitability. For each trade, we know 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.³¹ 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 8 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% of users with negative PnL into positive territory—roughly one in five losers would have ended the

³¹The 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¢.

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* the top 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 5 (Table 5), and is the realized-PnL counterpart of the maker–taker wealth transfer we examine in detail in Section 6.3.

Taken together, the decomposition shows that the bid-ask spread is a meaningful cost 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.³²

6.3. Maker-taker excess hit rate

Section 5 shows that users with a higher maker volume share are less likely to lose, and Section 6.2 shows 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. Hence, *why* does maker volume predict 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 act on through limit orders? The excess hit rate, defined as $o_k - p_k$, is constructed from the realized outcome and the traded price itself, so spread 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

³²Tables H16–H20 of the Internet Appendix 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.

driven entirely by the spread, 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 5 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 the prices at which 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.³³ 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

³³This pattern is the empirical signature predicted by competitive market-making models in the spirit of [Glosten and Milgrom \(1985\)](#) and [Kyle \(1985\)](#): in markets with deep order flow and strong inter-maker competition, informational rents earned by liquidity providers are compressed toward the rate that just covers the inventory cost of providing liquidity, while in less competitive or thinner markets makers can sustain a larger informational wedge. Recovering this special-case prediction in the data lends external validity to our cross-category interpretation.

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 5 and 6.2: 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 liquidity premium, makers earn a forecasting premium that is concentrated in the categories where prices are least efficient and mispricings are hardest for takers to exploit.

7. A Simulation Benchmark for Profit Concentration

The empirical results in Section 4 document both price efficiency and extreme profit concentration on Polymarket. We next examine how these facts can arise jointly. We simulate a parsimonious repeated binary-market economy in which traders differ in initial wealth, signal precision, perceived precision, and propensity to provide liquidity. Prices are generated from a public log-odds signal that is unbiased on average but contains small residual mispricings; traders decide whether to participate in each market and either take or provide liquidity. The model is not intended to structurally estimate what we find on Polymarket. It provides a benchmark for understanding which factors are sufficient to generate the distributional facts documented above. The exercise is organized around the question: can a market with unbiased aggregate prices nevertheless generate extreme concentration of trading gains and widespread losses among liquidity takers? The full

model, primitives, beliefs, participation, demand, and the maker-taker execution mechanism, is described in Appendix A. The simulation reported below uses $N = 100,000$ traders and $M = 50,000$ markets. The exercise is illustrative rather than structural: parameter values are set to highlight the mechanisms at play, not calibrated to match empirical moments.

Table 9 reports the cross-scenario summary alongside the empirical facts from Section 4. The exercise is anchored on a *Simulation baseline*, an economy with no information heterogeneity ($\tau_i = \tau$), no overconfidence ($\hat{\tau}_i = \tau_i$), and a homogeneous maker propensity ($\mu_i = \mu$); a constant half-spread is in place. The next three rows each turn on *one* channel on top of this baseline: *+ Skilled traders* reintroduces dispersion in the true signal precision τ_i ; *+ Maker specialization* reintroduces dispersion in the maker propensity μ_i ; *+ Overconfidence* reintroduces the perceived-precision wedge $\hat{\tau}_i \neq \tau_i$. Reading these rows against the simulation baseline isolates the marginal contribution of each channel. Row *+ All three* turns the three channels on simultaneously, with all primitives drawn independently across traders. Alongside the realized PnL, every specification carries a parallel reporting accumulator: the PnL each trader would have realized within the same simulation run had every one of their executed shares cleared at the market midpoint instead of at the bid or ask. This is a same-trade ex-post adjustment identical in spirit to the empirical spread decomposition in Section 6.2; the midpoint PnL does not feed back into wealth, participation, or any simulation dynamic, so for each row the trades, wealth paths, and execution quantities are unchanged from the realized economy. The exercise targets qualitative matches with the empirical pattern rather than numerical matches.

There are three takeaways. First, even the *Simulation baseline*, which switches off all three behavioral and informational channels, generates non-trivial winner concentration: the top 1% of winners capture 30.2% of gains and the top 5% capture 58.2%. With identical primitives and unbiased prices, dispersion in initial wealth and idiosyncratic luck on each trader’s signal sequence

are sufficient to push the gain distribution into a heavy-tailed regime. Some winner concentration is therefore mechanically generated by the trading environment itself, before any heterogeneity in skill, market-making, or behavioral biases is introduced. Second, on a stand-alone basis, skill heterogeneity and maker specialization each pull the top-share gain concentration further above the simulation baseline, while overconfidence has only a modest effect on its own: a perceived-precision wedge in the absence of true skill heterogeneity merely scales position sizes uniformly without selecting who wins. Third, the three channels are complementary. Turning them on together (*+ All three*) produces concentration moments that exceed the sum of the stand-alone deltas, indicating that the channels reinforce one another rather than acting additively. The rightmost column reports the share of losers whose PnL would have been positive had every one of their executed shares cleared at the market midpoint; this share is highest under *+ Maker specialization* and *+ All three*, where a non-trivial subset of losing traders are marginal in the sense that their losses are explained by spread costs rather than by being on the wrong side of the price.

Figures I4 and I5 in the Internet Appendix provide panel grids for the volume-weighted EHR and the maker share by PnL percentile across specifications. The simulated EHR gradient has the same sign as the empirical gradient and steepens once skill heterogeneity is introduced. The simulated maker-share profile is flat in the rows that hold μ_i constant and becomes monotone in the PnL rank under *+ Maker specialization* and *+ All three*.

The benchmark has clear limitations, the most important of which is its inability to reproduce the empirical loss-side concentration. In the data, the bottom 1% of losers bear 68.1% of total losses and the bottom 5% bear 86.1%; in the all-ingredients specification, the simulated counterparts are only 28.3% and 55.1%. The channels we vary generate concentration on the winner side but barely move the loser side. The joint distribution of trader primitives combined with the loss-capping mechanics of the model can explain this result. Empirically, the most concentrated losers are plausibly traders

who are simultaneously *overconfident*, *low-skill*, and *high-wealth*, a combination that produces large absolute losses. In our independent-primitive draws, this combination is rare; and even when it occurs, the loss-capping property of the model truncates the trader’s downside at their initial wealth (and exits them via the participation rule once their wealth dips below $W_{\min}$), while a winning trader’s gains are uncapped and compound through the wealth-dependent participation channel. The result is an asymmetric simulation: extreme winners can keep accumulating, but extreme losers disappear too quickly to dominate the loss tail. Closing this gap would require introducing a positive correlation between overconfidence, low skill, and high wealth, and either weakening the exit threshold or allowing wealth replenishment so that committed losers can keep losing.

8. Trading Behavior for the Most Profitable Users

Finally, we conduct a descriptive analysis of the trading behavior of the 100 most successful users in our sample, who together account for 27.7 percent of aggregate winners’ profits in our sample. There has been recent commentary in the media, by practitioners, and among academics speculating as to the role of insider information in trading profits. Such evidence frequently reports a suspicious pattern of trading that is defined *ex post* and comments on the magnitude of profits without contextualizing the magnitudes in light of the total size and scope of Polymarket.³⁴ Our results on user EHRs, performance persistence, and market making suggest that there are potentially a number of different drivers of the profits of the most profitable traders. Our analysis so far has not examined the role of insider information specifically. While we do not think it is possible to conclusively prove or disprove the extent of such behavior, we are interested in seeing whether the trading patterns among the most profitable users could plausibly be coming from trading in markets where insider information would be valuable.

³⁴See for example, <https://www.nytimes.com/2026/05/13/technology/polymarket-insider-trading.html>.

In Section J of the Internet Appendix, we list the top 100 winners along with the markets in which they gained or lost the most. Two patterns within the top 100 deserve emphasis, and both bear on the question of what drives their profits. First, individual top earners are not all broadly diversified across markets: most have their PnL concentrated in a small number of markets, and in many cases the bulk of a user’s terminal PnL is attributable to a single market or a tight cluster of related markets, as illustrated by the “largest market share of total volume” column in Figure J6 and Tables J22–J23 of the Internet Appendix. The largest source of profits is sports betting, accounting for 81% of total gains, spread across a variety of sports and teams with no single one dominating. The second source of profits comes from the political category, with the largest bets on “Will Donald Trump win the 2024 US Presidential Election?,” “Will Donald Trump win the popular vote in the 2024 Presidential Election?,” and “Kamala Harris wins the popular vote?” accounting for 15% of total gains. The remaining 4% of gains are distributed across a variety of other categories. Second, a meaningful fraction of the top 100 are pure liquidity providers: they trade extensively, post limit orders across many markets, and earn a small amount on each trade rather than placing a few large directional bets, showing that liquidity provision can itself be very profitable. In order for “insider” trading to explain the performance of the non-market-making accounts an insider source would have to span dozens of sports leagues and have material non-public knowledge of the voting intentions of tens of millions of Americans for it to explain the bulk of the observed gains.

While we do not claim that there is *no* insider trading on Polymarket, the most successful traders do not seem to have traded in ways that, at a high level, are consistent with a large degree of insider trading. Rather, to the extent that such traders have some form of skill, our previous results suggest that it is more likely to be a combination of successful liquidity provision and better forecasting or processing of publicly available information. Naturally, it is not possible to say with certainty that none of the traders in this sample possessed insider information, but we urge caution when

concluding that insider trading is pervasive on this platform, at least from the perspective of the aggregate volume of profits.

9. Conclusion and Implications for Future Research

Our analysis of over 2.4 million Polymarket users reveals five findings. First, profits are highly concentrated in a small fraction of users, while the majority lose money, consistent with a near-zero-sum market. Second, prices are relatively efficient and informative on aggregate, though low-volume markets and the Tech, Culture, and Weather categories exhibit systematic deviations from perfect calibration. Third, per-trade forecasting accuracy declines sharply through the PnL distribution, suggesting that profitability is associated with the ability to identify mispriced contracts, though monthly performance is only weakly persistent. Fourth, liquidity provision is the strongest cross-sectional predictor of avoiding losses. Fifth, a descriptive analysis of the most successful users shows that their gains span many sports leagues and a handful of 2024 US election markets, with a meaningful share from pure liquidity provision; the breadth of these patterns argues against insider trading as the dominant source of their profits. A parsimonious simulation benchmark shows that heterogeneity in skill, liquidity-provision propensity, and overconfidence are jointly sufficient to reproduce the empirical concentration and forecasting-accuracy patterns. For investors, prediction markets should be treated as other highly competitive financial markets: prices are on average fair, market-level PnL is zero-sum, and the probability of loss is higher for liquidity takers.

Several questions remain open. The behavioral mechanisms underlying low market-making activity, such as overconfidence, probability misperception, or a preference for lottery-like payoffs, deserve more empirical investigation using trader-level data. As prediction markets grow in scale and visibility, understanding their interaction with traditional financial markets also becomes 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. We provide the data used in this paper to facilitate future research.

A. Simulation Benchmark: Model Details

This appendix lays out the simulation benchmark referenced in Section 7. We simulate M independent binary markets and N traders. Each market resolves once and is matched against a single bid-ask quote per trade; traders enter or exit endogenously based on their accumulated wealth. Holding the overall pricing technology fixed, the simulation traces how the cross-section of trading outcomes responds to differences in skill, overconfidence, market-making behavior, and survival. The exercise links directly to the empirical findings of Sections 4.1–6.3: it produces simulated counterparts of the concentration of gains and losses, the EHR gradient by PnL percentile, the maker-taker asymmetry, and the spread decomposition.

A.1. Market primitives

There are M binary markets indexed by $m = 1, \dots, M$. Each market has a latent true probability and a realized outcome,

$$\theta_m \sim \text{Uniform}(0, 1), \quad Y_m \sim \text{Bernoulli}(\theta_m). \quad (2)$$

We work in log-odds space, $z_m = \log(\theta_m/(1 - \theta_m))$. The public log-odds quote is unbiased on average but contains residual pricing noise,

$$\ell_m = z_m + u_m, \quad u_m \sim N(0, \sigma_p^2), \quad p_m = \Lambda(\ell_m), \quad (3)$$

where $\Lambda(x) = (1 + e^{-x})^{-1}$. This structure lets simulated prices be unbiased on average in log-odds space while still leaving small exploitable deviations for better-informed traders. A category-specific noise term $\sigma_{p,c(m)}$ would let efficient categories such as Sports and Crypto have lower residual mispricing and less efficient categories such as Politics and Weather have larger residual mispricing;

we use a category-free σ_p in the benchmark for transparency.

A.2. Trader heterogeneity

There are N traders indexed by $i = 1, \dots, N$. Each trader is characterized by four primitive features: initial wealth, true signal precision, maker propensity, and overconfidence. The four features are drawn independently from their own (log-)normal distributions, with no common latent factor.

Initial wealth is

$$\log W_{i0} = \alpha_W + \varepsilon_i^W, \quad \varepsilon_i^W \sim N(0, \sigma_W^2). \quad (4)$$

True signal precision is

$$\log \tau_i = \alpha_\tau + \varepsilon_i^\tau, \quad \varepsilon_i^\tau \sim N(0, \sigma_\tau^2). \quad (5)$$

The baseline maker propensity is

$$\mu_i = \Lambda(\alpha_\mu + \varepsilon_i^\mu), \quad \varepsilon_i^\mu \sim N(0, \sigma_\mu^2). \quad (6)$$

Perceived signal precision adds an overconfidence wedge³⁵,

$$\log \hat{\tau}_i = \log \tau_i + b_i, \quad b_i = \alpha_b + \sigma_b \eta_i^b, \quad \eta_i^b \sim N(0, 1). \quad (7)$$

Independent draws across the four primitives mean that, in the benchmark, the high-wealth tail of the cross-section is not the same set of traders as the high-precision tail, the high-maker-propensity tail, or the most-overconfident tail. This independence assumption is restrictive: empirically, sophisticated traders are plausibly both better informed, more likely to provide liquidity, and *less* overconfident than unsophisticated traders. We adopt the independence assumption in this paper

³⁵The concept of an overconfidence wedge is similar to [Odean \(1998\)](#) and the heterogeneous information precision follows [Grossman and Stiglitz \(1980\)](#).

for transparency and leave a richer joint distribution of trader primitives for future work.

A.3. Signals and beliefs

Trader i receives a noisy signal about the true log-odds of market m ,

$$x_{im} = z_m + \eta_{im}, \quad \eta_{im} \sim N(0, 1/\tau_i). \quad (8)$$

The trader forms a subjective posterior log-odds by combining the private signal with the publicly observable log-odds quote ℓ_m ,

$$\tilde{z}_{im} = \lambda_i x_{im} + (1 - \lambda_i) \ell_m, \quad \lambda_i = \frac{\hat{\tau}_i}{\hat{\tau}_i + \tau_0}, \quad q_{im} = \Lambda(\tilde{z}_{im}). \quad (9)$$

This is the Bayesian posterior of a trader who treats x_{im} as a private signal with precision $\hat{\tau}_i$ and the public quote as a separate signal with precision τ_0 . With $\tau_0 = 1/\sigma_p^2$ the weights λ_i are exactly Bayesian-rational; deviations from this baseline through $\hat{\tau}_i \neq \tau_i$ capture overconfidence. Sophisticated traders ($\hat{\tau}_i \gg \tau_0$) put weight on their private signal; unsophisticated traders without overconfidence ($\hat{\tau}_i \ll \tau_0$) defer to the public price; overconfident unsophisticated traders place outsized weight on their noisy signal, which feeds into their position sizes through Equation (11).

A.4. Participation and survival

Trader i participates in market m with probability

$$\Pr(A_{im} = 1) = \Lambda(\alpha_A + \beta_{A,W} \log W_{it}). \quad (10)$$

Only active traders with positive wealth above the threshold $W_{\min}$ can participate; traders below the threshold become inactive.³⁶ Wealth-dependent participation creates a compounding channel: traders who win accumulate wealth and trade more often, while traders who lose become less active or exit. This is the simulation analog of the empirical pattern in Table 6 that lower performance positively predicts exit.

A.5. Trading demand

For participating traders, define the subjective edge $e_{im} = q_{im} - p_m$. Desired signed notional demand is

$$d_{im}^* = \kappa_i W_{it} \frac{q_{im} - p_m}{p_m(1 - p_m)}, \quad \kappa_i = \kappa_0 \hat{\tau}_i^\gamma, \quad (11)$$

and is capped as a fraction of wealth³⁷,

$$d_{im} = \text{med}(-\bar{\omega}W_{it}, d_{im}^*, \bar{\omega}W_{it}). \quad (12)$$

Positive d_{im} means the trader wants to buy the reference outcome; negative d_{im} means the trader wants to buy the complementary outcome. Overconfidence enters through $\hat{\tau}_i$ in κ_i : traders with noisy signals but inflated perceived precision take systematically larger positions than their information warrants.

A.6. Maker-taker execution

We use a constant-half-spread bid-ask mechanism rather than a strategic auction. The execution prices are

$$p_m^{ask} = \min(p_m + h_0, 1), \quad p_m^{bid} = \max(p_m - h_0, 0), \quad (13)$$

³⁶Blume and Easley (1992) explore the dynamics of wealth accumulation and its implications for the long run behavior of asset prices.

³⁷This demand is similar to the Kelly (1956) criterion for optimal bet sizing applied to binary bets.

and a participating trader is a maker with probability μ_i in each market. Within a market we run two independent crosses. Cross A matches takers wanting long-reference exposure ($d_{im} > 0$) against makers wanting short-reference exposure ($d_{im} < 0$) at p_m^{ask} ; cross B matches takers wanting short-reference exposure ($d_{im} < 0$) against makers wanting long-reference exposure ($d_{im} > 0$) at p_m^{bid} . Each notional demand is converted to desired signed reference shares at the prevailing mid: a reference-outcome buy ($d_{im} > 0$) is converted at d_{im}/p_m , while a complement-outcome buy ($d_{im} < 0$) is converted at $d_{im}/(1 - p_m)$. Within each cross, both sides are pro-rated when imbalanced. Makers earn the spread; the trading PnL is approximately zero-sum across all participants in each market (no fees in this benchmark).

A.7. PnL and wealth dynamics

Per-share PnL is $\pi_{im} = Y_m - p_{im}^{exec}$ for buyers of the reference outcome and $\pi_{im} = p_{im}^{exec} - Y_m$ for buyers of the complementary outcome; total PnL is per-share PnL multiplied by the absolute share count, wealth evolves as $W_{i,t+1} = W_{it} + \Pi_{im}$, and traders exit when $W_{it} \leq W_{\min}$.

A.8. Simulation exercises

The exercise is anchored on a *Simulation baseline* that switches off all three behavioral and informational channels at once. The next three rows each turn on *one* channel on top of this baseline in isolation; the final row then turns all three channels on jointly. The specifications are:

- **Simulation baseline:** no information heterogeneity ($\tau_i = \tau$ for all i), no overconfidence ($\hat{\tau}_i = \tau_i$), and homogeneous maker propensity ($\mu_i = \mu$ for all i). The constant half-spread h_0 is in place. All trader-level dispersion in this specification comes from initial wealth and idiosyncratic participation shocks.
- **+ Skilled traders:** simulation baseline plus cross-sectional dispersion in the true signal

precision τ_i only.

- + **Maker specialization:** simulation baseline plus cross-sectional dispersion in the maker propensity μ_i only.
- + **Overconfidence:** simulation baseline plus the perceived-precision wedge ($\hat{\tau}_i \neq \tau_i$ as in Equation (7)) only. Skill and maker propensity remain homogeneous.
- + **All three:** simulation baseline plus dispersion in τ_i , μ_i , and the overconfidence wedge b_i simultaneously, all drawn independently. This row corresponds to the full benchmark with no scenario overrides.

For each specification, we additionally track a parallel *midpoint* PnL accumulator: per trade, in addition to the realized PnL “ $s_{i,m}(Y_m - q_{i,m})$ ” that updates wealth, we accumulate a hypothetical PnL “ $s_{i,m}(Y_m - p_m)$ ” computed at the market mid. The midpoint PnL is identical in spirit to the empirical spread decomposition in Section 6.2: it asks what each trader’s outcome would have been if their same executed shares had cleared at the midpoint instead of at the bid or ask. It does not feed back into wealth, participation, or any simulation dynamic, so within a row it shares the same trades, wealth path, and execution quantities as the realized economy. The rightmost column of Table 9 reports, for each specification, the fraction of users with negative realized cumulative PnL whose midpoint PnL is positive.

Each specification differs from the simulation baseline only through configuration overrides; all other primitives, including the random seed, are held fixed. Calibration is loose: we hand-tune a small set of parameters so that the all-ingredients specification produces broad concentration magnitudes consistent with Section 4.1 and Table 5, not exact moments. The simulation reported in Section 7 uses $N = 100,000$ traders and $M = 50,000$ markets, which gives every trader many trading rounds and keeps the number of users in the top 0.1% slice large enough to estimate top-share

statistics.

References

- 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, Working paper.
- 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., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, 2014, The cross-section of speculator skill: Evidence from day trading, *Journal of Financial Markets* 18, 1–24.
- Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *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.
- Blume, Lawrence, and David Easley, 1992, Evolution and market behavior, *Journal of Economic theory* 58, 9–40.
- 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.
- Chernov, Mikhail, Vadim Elenev, and Dongho Song, 2026, The comovement of voter preferences: Insights from us presidential election prediction markets beyond polls, *Journal of Econometrics* p. 106192.
- Cookson, J. Anthony, 2018, When saving is gambling, *Journal of Financial Economics* 129, 24–45.
- , Marina Niessner, and Christoph Schiller, 2026, Can social media inform corporate decisions? Evidence from merger withdrawals, *Journal of Finance* 81, 91–142.
- Cziraki, Peter, and Jasmin Gider, 2021, The dollar profits to insider trading, *Review of Finance* 25, 1547–1580.

- Della Vedova, Joshua, 2026, Who profits from prediction markets? Execution, not information, *Working paper*.
- Dewally, Michaël, Louis H. Ederington, and Chitru S. Fernando, 2013, Determinants of trader profits in commodity futures markets, *Review of Financial Studies* 26, 2648–2683.
- Diercks, Anthony M., Jared Dean Katz, and Jonathan H. Wright, 2026, Kalshi and the rise of macro markets, *Working paper*.
- Dorn, Anne Jones, Daniel Dorn, and Paul Sengmueller, 2015, Trading as gambling, *Management Science* 61, 2376–2393.
- Forsythe, Robert, Forrest Nelson, George R. Neumann, and Jack Wright, 1992, Anatomy of an experimental political stock market, *American Economic Review* 82, 1142–1161.
- Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Gomez Cram, Roberto, Yunhan Guo, Theis Ingerslev Jensen, and Howard Kung, 2025, Financial prediction markets: A new measure of earnings expectations, *Working paper*.
- , 2026, Prediction market accuracy: Crowd wisdom or informed minority?, *Working paper*.
- Grégoire, Vincent, and Charles Martineau, 2022, How is earnings news transmitted to stock prices?, *Journal of Accounting Research* 60, 261–297.
- Grossman, Sanford J., and Joseph E Stiglitz, 1980, On the impossibility of informationally efficient markets, *The American economic review* 70, 393–408.
- Hayek, Friedrich A., 1945, The use of knowledge in society, *American Economic Review* 35, 519–530.
- Hu, Jianfeng, Antonia Kirilova, Seongkyu (Gilbert) Park, and Doojin Ryu, 2024, Who profits from trading options?, *Management Science* 70, 4742–4761.
- Kelly, John L, 1956, A new interpretation of information rate, *the bell system technical journal* 35, 917–926.
- Kumar, Alok, 2009, Who gambles in the stock market?, *Journal of Finance* 64, 1889–1933.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

- 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.
- Ng, Hunter, Lin Peng, Yubo Tao, and Dexin Zhou, 2026, Price discovery and trading in modern prediction markets, *Working paper*.
- Odean, Terrance, 1998, Volume, volatility, price, and profit when all traders are above average, *Journal of Finance* 53, 1887–1934.
- , 1999, Do investors trade too much?, *American Economic Review* 89, 1279–1298.
- 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 Martin Walther, 2025, Exploring decentralized prediction markets: Accuracy, skill, and bias on Polymarket, *Working paper*.
- Sirolly, Allen, Hongyao Ma, Yash Kanoria, and Rajiv Sethi, 2025, Network-based detection of wash trading, *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, 2013, Prediction markets for economic forecasting, in Graham Elliott, and Allan Timmermann, ed.: *Handbook of Economic Forecasting* vol. 2 . pp. 657–687 (Elsevier).
- 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, 2026, 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.

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

Figure 1: Polymarket Growth: Users, Volume, and New Markets

This figure summarizes Polymarket’s platform growth. Panel (a) plots monthly new users (bars) with cumulative users (solid line). Panel (b) plots monthly trading volume in USDC, stacked by category, with cumulative volume (solid line). Panel (c) plots the number of new events and markets created each month on a log scale. Panels (a) and (b) cover the trade sample, which runs from November 11, 2022 to March 29, 2026. Panel (c) covers all markets opened on the platform from Polymarket’s launch in mid-2020 through March 29, 2026.




Figure 2: Calibration of Polymarket Prices

This figure plots the calibration of Polymarket binary market prices at different time horizons before resolution for all markets and 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 3: Concentration of Gains and Losses Across Polymarket Users

This figure summarizes the concentration of trading gains and losses among Polymarket users. Panel (a) plots the cumulative concentration as a snapshot at the end of the sample: 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, and the red bar reports the share of total losses borne by users in the corresponding bottom group. Panel (b) plots the shares over time at a monthly frequency, with one line per percentile bucket (0.1%, 1%, 5%) for gains and losses. Panel (c) plots the fraction of users with positive PnL at each month-end. The sample period is from November 11, 2022 to March 29, 2026.




Figure 4: Prediction Accuracy of Winners and Losers

This figure plots the volume-weighted excess hit rate $\text{EHR}_{\text{vw},i} = \sum_k q_k (o_k - p_k) / \sum_k q_k$ by PnL percentile. Users are sorted in ascending PnL order and assigned to 200 half-percentile bins; the x -axis runs from percentile 0 (least profitable, large negative PnL) on the left to percentile 100 (most profitable) on the right. Within each bin we plot the mean (solid blue, with a 95% confidence band) and the median (dot-dashed black). The horizontal dashed line is the sample-mean EHR_{vw} across all users with at least one resolved trade. 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 if the contract purchased in trade k (after buy-side normalization) resolves in-the-money, otherwise equals 0, 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. EHR is defined only over resolved trades, so users with no resolved trades during the sample period are excluded from the figure; the PnL ranking on the x -axis remains over all users. The sample period is from November 11, 2022 to March 29, 2026.




Figure 5: Excess Hit Rate Differences Between Makers and Takers

This figure plots the difference in excess hit rate (Equation 1) 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 thin dashed horizontal line at zero marks the null of no maker-taker gap; the thicker dashed horizontal line marks the trade-fraction-weighted average gap across all price bins in the panel. The gray filled area plots the fraction of total trades at each price bin on the right y-axis. The x-axis spans the upper half of the buy-price domain (\$0.50 to \$1.00); the secondary axis at the top shows the equivalent sell price (\$0.50 down to \$0.00) under buy-side normalization, so each bin corresponds to both directions of trade. The vertical line at buy price \$0.96 (equivalently sell price \$0.04) marks the small-tick boundary, beyond which maker quotes cross the spread at sub-penny ticks. EHR is defined only over resolved trades. 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 Polymarket users in our sample. *Trades per User* is the total number of executed trades. *Markets per User* is the number of distinct binary markets traded. *Volume per User (USDC)* is total notional volume in USDC, applying the matched-notional convention in which one share traded counts as one USDC of volume. *Fraction Extreme Price* is the share of a user's trades executed at prices below 10¢ or above 90¢. *PnL per User (USDC)* is the final mark-to-market profit and loss in USDC, 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). All monetary values are in USDC. Statistics reported are the mean, standard deviation, minimum, 5th, 25th, 50th, 75th, and 95th percentiles, and maximum. Panel A reports statistics for all users, Panel B for users with more than 100 trades, and Panel C for users with more than 1,000 trades. The sample period is from November 11, 2022 to March 29, 2026.

Metric	Mean	Std. Dev.	Min	5th Pct.	25th Pct.	Median	75th Pct.	95th Pct.	Max
Panel A: All users ($N = 2,480,104$)									
Trades per User	474	21,450	1	1	9	28	74	546	13,587,640
Markets per User	54	495	1	1	3	10	26	127	130,387
Volume per User (USDC)	54,143	1,699,528	0	20	474	2,343	11,619	97,793	604,671,119
Fraction Extreme Price	54.9%	39.7%	0.0%	0.0%	14.0%	58.3%	100.0%	100.0%	100.0%
PnL per User (USDC)	0	32,376	-10,022,403	-920	-32	-2	1	252	22,029,973
Fraction Maker Volume	17.9%	27.2%	0.0%	0.0%	0.0%	0.0%	39.5%	75.4%	100.0%
Panel B: Users with more than 100 trades ($N = 478,096$)									
Trades per User	2,349	48,809	101	107	137	221	576	4,552	13,587,640
Markets per User	231	1,110	1	9	31	61	132	800	130,387
Volume per User (USDC)	234,116	3,863,332	13	1,129	4,557	15,411	50,156	447,591	604,671,119
Fraction Extreme Price	48.7%	35.2%	0.0%	1.7%	17.1%	39.7%	87.3%	100.0%	100.0%
PnL per User (USDC)	181	73,609	-10,022,403	-2,688	-147	-14	5	1,188	22,029,973
Fraction Maker Volume	24.2%	27.9%	0.0%	0.0%	0.1%	9.3%	47.3%	81.0%	100.0%
Panel C: Users with more than 1,000 trades ($N = 78,366$)									
Trades per User	12,957	119,998	1,001	1,070	1,473	2,509	5,673	29,905	13,587,640
Markets per User	1,023	2,594	1	45	189	433	971	3,583	130,387
Volume per User (USDC)	1,150,361	9,464,494	152	6,920	24,393	78,459	329,973	3,111,043	604,671,119
Fraction Extreme Price	31.5%	29.1%	0.0%	1.3%	9.5%	20.6%	45.7%	96.7%	100.0%
PnL per User (USDC)	2,627	179,118	-10,022,403	-11,871	-916	-140	124	12,686	22,029,973
Fraction Maker Volume	31.9%	31.2%	0.0%	0.0%	1.6%	23.5%	55.4%	91.9%	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* is in millions, *Volume* is in millions of USDC, and both are also reported as a percentage of the 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). *Trades/Markets* and *Volume/Markets* are per-market averages, reported in raw counts and USDC respectively. 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. The sample period is from November 11, 2022 to March 29, 2026.

Category	Events	Markets	Trades (M)	Trades (%)	Volume (M USDC)	Volume (%)	Users	Users (%)	Trades/Markets	Volume/Markets
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 ($\text{PnL} < 0$) and winners ($\text{PnL} > 0$), 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	614,883	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 (M USDC)	67,135.4	25,477.0	21,114.4	12,762.0	3,473.9	3,107.5	806.9	393.7
Frac. Winners	30.9%	31.5%	35.5%	42.5%	38.6%	46.6%	50.3%	52.6%
Mean PnL by Percentile – Losers ($\text{PnL} < 0$)								
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 ($\text{PnL} > 0$)								
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: Probit Regression: Determinants of Negative Performance.

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*** (0.001)	−0.016*** (0.001)	−0.006*** (0.001)	−0.003*** (0.001)	0.007*** (0.002)	0.012*** (0.002)	−0.129*** (0.006)	−0.128*** (0.006)
Frac Maker Volume		−0.354*** (0.001)	−0.349*** (0.001)	−0.344*** (0.001)	−0.327*** (0.002)	−0.324*** (0.002)	−0.147*** (0.006)	−0.142*** (0.006)
Log N Trades		−0.001*** (0.000)	0.004*** (0.000)	0.017*** (0.000)	0.027*** (0.001)	0.029*** (0.001)	−0.004** (0.002)	−0.006*** (0.002)
Log Total Volume		0.020*** (0.000)	0.018*** (0.000)	0.020*** (0.000)	−0.021*** (0.000)	−0.021*** (0.001)	−0.043*** (0.001)	−0.041*** (0.001)
Category HHI			0.081*** (0.001)	−0.055*** (0.002)	−0.010*** (0.003)	−0.067*** (0.004)	−0.131*** (0.007)	−0.143*** (0.010)
Counterparty HHI			−0.005*** (0.002)	0.036*** (0.002)	0.180*** (0.007)	0.206*** (0.007)	0.500*** (0.050)	0.494*** (0.049)
Traded Crypto				−0.074*** (0.001)		−0.016*** (0.002)		0.057*** (0.006)
Traded Finance				−0.033*** (0.001)		−0.018*** (0.002)		−0.010* (0.005)
Traded Politics				−0.056*** (0.001)		0.000 (0.002)		0.041*** (0.005)
Traded Tech				−0.033*** (0.001)		−0.023*** (0.002)		−0.026*** (0.005)
Traded Culture				0.002*** (0.001)		−0.006*** (0.002)		−0.024*** (0.005)
Traded Weather				−0.036*** (0.001)		−0.021*** (0.002)		−0.008 (0.005)
Observations	2,480,104	2,480,104	2,480,104	2,480,104	478,096	478,096	78,366	78,366
Pseudo R^2	0.000	0.034	0.036	0.043	0.041	0.042	0.054	0.056

Table 5: PnL Group Characteristics

This table reports trading characteristics for users sorted into four non-overlapping PnL groups: top 0.1%, next 0.9%, next 4%, and bottom 95%. Panel A reports cross-sectional means of behavioral characteristics; Panel B reports risk-adjusted performance metrics computed at the user level from the monthly PnL series, followed by the cross-sectional mean of each user's excess hit rate. Columns (5)–(7) report the difference between each top group's statistic and the bottom 95% statistic. Significance stars in the difference columns (5)–(7) are computed from Welch's two-sample t -test for all rows except the two median rows in Panel B (*Median Ann. PnL-to-Volatility* and *Median Ann. PnL-to-Downside-Volatility*), which use the Mann-Whitney U test: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. See text for variable definitions and the user-month inclusion rule. The sample period is from November 11, 2022 to March 29, 2026.

	PnL Group Mean				Difference		
	Top 0.1% (1)	Next 0.9% (2)	Next 4% (3)	Bot. 95% (4)	0.1%–95% (5)	0.9%–95% (6)	4%–95% (7)
<i>Panel A: Behavioral Characteristics</i>							
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***
Avg N Trades	150,651	8,393	1,166	222	150,429***	8,171***	944***
<i>Panel B: Risk-Adjusted Performance</i>							
Mean Monthly PnL (USDC)	54,606.82	2,474.24	311.27	-104.62	54,711.44***	2,578.87***	415.90***
SD Monthly PnL (USDC)	104,098.08	5,106.82	640.34	220.39	103,877.69***	4,886.44***	419.95***
Max Drawdown (USDC)	35,212.20	1,897.74	228.57	366.02	34,846.18***	1,531.72***	-137.45***
Median Ann. PnL-to-Volatility	1.732	1.469	1.461	-1.150	2.882***	2.618***	2.610***
Median Ann. PnL-to-Downside-Volatility	12.128	20.445	12.191	-1.361	13.489***	21.806***	13.552***
Excess Hit Rate	0.0974	0.1317	0.1027	-0.0265	0.1238***	0.1582***	0.1292***

Table 6: Transition Matrix of PnL Groups

This table reports transition probabilities across monthly-performance groups over 1-month and 3-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 eleven 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, using the breakpoints $-100k$, $-10k$, $-1k$, -100 , -10 , 10 , 100 , $1k$, $10k$, and $100k$ (in dollars). The eleven bucket columns report transition frequencies *conditional on trading* (each row's bucket-to-bucket cells sum to 100% and are shaded in blue). The “ > 0 ” column reports the conditional probability that the trader's target-month PnL change is strictly positive. Significance stars come from a two-sided exact-binomial test against the null hypothesis that the probability of profit equals 0.5 (i.e., that profits and losses are equally likely): * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Cells are shaded only when the test rejects the null at the 10% level: green for probabilities above 0.5, red for probabilities below 0.5. The “No Trade” column reports the unconditional share of users in each starting group with no positions in the target month, shaded in blue. Long-horizon results (6-month and 12-month) are reported in Internet Appendix Table F10. The sample period is from November 11, 2022 to March 29, 2026.

Panel A: 1-Month Horizon														
Starting Group	N	< -100k	[-100k, -10k)	[-10k, -1k)	[-1k, -100)	[-100, -10)	[-10, 10]	(10, 100]	(100, 1k]	(1k, 10k]	(10k, 100k]	> 100k	> 0	No Trade
< -100k	885	18.1%	21.3%	9.1%	2.3%	0.4%	0.5%	0.4%	1.4%	8.2%	21.5%	16.8%	48.7%	36.8%
[-100k, -10k)	9,470	2.0%	20.4%	25.6%	8.2%	1.9%	4.6%	1.5%	5.4%	14.0%	14.2%	2.0%	40.1%***	34.3%
[-10k, -1k)	67,617	0.1%	3.7%	22.9%	20.2%	9.0%	12.2%	5.2%	10.8%	13.3%	2.6%	0.1%	39.4%***	41.9%
[-1k, -100)	197,668	0.0%	0.3%	6.4%	25.7%	19.7%	17.8%	11.3%	13.8%	4.5%	0.4%	0.0%	39.5%***	45.4%
[-100, -10)	437,629	0.0%	0.0%	0.9%	7.1%	24.5%	46.9%	15.4%	4.5%	0.7%	0.1%	0.0%	44.2%***	42.7%
[-10, 10]	3,260,841	0.0%	0.0%	0.2%	0.8%	5.1%	89.3%	3.8%	0.6%	0.1%	0.0%	0.0%	57.4%***	36.2%
(10, 100]	307,046	0.0%	0.0%	1.2%	8.3%	23.0%	39.0%	21.7%	6.0%	0.8%	0.1%	0.0%	51.9%***	31.0%
(100, 1k]	132,379	0.0%	0.4%	7.1%	23.4%	12.9%	15.4%	12.4%	21.0%	6.8%	0.5%	0.0%	50.0%	31.2%
(1k, 10k]	49,780	0.1%	3.9%	20.8%	12.9%	4.1%	10.6%	4.0%	12.7%	25.3%	5.5%	0.1%	54.3%***	23.8%
(10k, 100k]	10,481	1.9%	17.0%	14.1%	3.7%	1.7%	3.6%	1.5%	3.8%	19.2%	30.4%	3.1%	60.5%***	20.0%
> 100k	1,065	17.8%	15.4%	6.3%	1.2%	0.1%	0.2%	0.2%	0.3%	5.7%	22.8%	29.8%	59.0%***	16.7%
Panel B: 3-Month Horizon														
Starting Group	N	< -100k	[-100k, -10k)	[-10k, -1k)	[-1k, -100)	[-100, -10)	[-10, 10]	(10, 100]	(100, 1k]	(1k, 10k]	(10k, 100k]	> 100k	> 0	No Trade
< -100k	679	14.6%	18.4%	17.3%	3.1%	2.4%	0.7%	0.3%	3.4%	6.5%	21.4%	11.9%	43.9%**	56.7%
[-100k, -10k)	7,169	2.3%	16.3%	22.4%	9.9%	2.8%	4.3%	2.0%	8.0%	17.0%	12.5%	2.5%	44.8%***	52.4%
[-10k, -1k)	45,077	0.2%	3.5%	19.1%	18.4%	8.7%	15.9%	7.0%	11.5%	12.4%	3.1%	0.2%	44.3%***	52.5%
[-1k, -100)	134,581	0.0%	0.6%	6.2%	20.7%	17.6%	24.5%	12.5%	12.4%	4.8%	0.6%	0.0%	44.7%***	58.2%
[-100, -10)	314,294	0.0%	0.1%	1.2%	6.6%	20.0%	53.0%	13.6%	4.4%	1.0%	0.1%	0.0%	47.6%***	57.7%
[-10, 10]	2,768,522	0.0%	0.0%	0.4%	1.0%	4.7%	88.9%	3.8%	0.8%	0.3%	0.0%	0.0%	60.2%***	52.5%
(10, 100]	231,008	0.0%	0.1%	1.2%	6.7%	17.1%	51.0%	17.5%	5.4%	1.1%	0.1%	0.0%	54.1%***	49.3%
(100, 1k]	94,689	0.0%	0.6%	6.5%	17.1%	13.7%	23.4%	12.7%	17.7%	7.4%	0.8%	0.0%	52.9%***	47.7%
(1k, 10k]	36,798	0.2%	3.7%	15.7%	14.0%	6.5%	11.8%	5.4%	13.7%	22.6%	6.1%	0.2%	55.6%***	38.7%
(10k, 100k]	7,567	2.0%	13.5%	15.1%	6.3%	2.0%	4.0%	1.7%	5.5%	19.0%	27.0%	4.0%	60.1%***	33.2%
> 100k	794	11.3%	14.4%	6.1%	4.0%	0.2%	0.2%	0.4%	1.7%	9.8%	29.9%	22.1%	64.1%***	39.7%

Table 7: Risk-Adjusted PnL Transition Matrix — Sharpe Analog

This table replicates Table 6 after replacing the dollar-PnL bucketing with an annualized risk-adjusted ratio (a Sharpe analog computed at the daily frequency). In each month, users with at least 10 active days are sorted into nine groups based on the annualized $\sqrt{365} \cdot \overline{\Delta \text{PnL}_d} / \sigma(\Delta \text{PnL}_d)$ ratio of their daily mark-to-market PnL changes during the month, with cutoffs at ± 10 , ± 5 , ± 3 , and ± 1 (the merged middle bucket $[-1, 1]$ collapses what would otherwise be three narrow groups around the breakeven point). Bucket cells report the percentage of users in each starting group whose ratio falls into each group at the specified horizon, conditional on a valid target-month observation. The “ > 0 ” column reports the conditional probability of a strictly positive target-month ratio, with stars from a two-sided exact-binomial test against $H_0: p = 0.5$ (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). The No Trade column reports the unconditional rate of users with no qualifying target-month observation (no trades or insufficient activity). The Sortino-analog version (which replaces the denominator with the downside semi-deviation) and the 6- and 12-month long-horizon variants are reported in Internet Appendix Section G. The sample period is from November 11, 2022 to March 29, 2026.

Panel A: 1-Month Horizon

Starting Group	N	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	51,035	5.7%	19.5%	17.7%	18.3%	20.4%	9.1%	5.6%	3.2%	0.5%	28.0%***	45.6%
[-10, -5)	503,828	1.1%	12.7%	18.6%	19.7%	24.6%	11.3%	7.7%	4.0%	0.5%	35.0%***	44.6%
[-5, -3)	749,262	0.5%	9.3%	18.1%	17.9%	27.4%	12.1%	9.8%	4.4%	0.4%	39.9%***	41.5%
[-3, -1)	754,860	0.4%	7.2%	13.0%	17.6%	35.2%	14.4%	7.8%	4.1%	0.3%	45.0%***	24.8%
[-1, 1]	1,447,976	0.2%	4.0%	7.9%	13.6%	54.9%	11.2%	5.3%	2.9%	0.2%	45.8%***	25.4%
(1, 3]	576,792	0.3%	6.2%	12.7%	18.4%	29.9%	14.6%	10.8%	6.5%	0.6%	45.7%***	25.5%
(3, 5]	436,182	0.3%	6.5%	15.8%	14.2%	22.0%	15.7%	15.3%	9.2%	1.0%	52.4%***	33.4%
(5, 10]	295,893	0.4%	6.6%	12.7%	12.6%	17.1%	16.1%	16.3%	15.3%	2.9%	59.9%***	36.2%
> 10	43,726	0.6%	7.5%	11.5%	10.2%	9.9%	9.3%	10.8%	22.9%	17.2%	65.3%***	33.1%

Panel B: 3-Month Horizon

Starting Group	N	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	35,545	1.8%	10.3%	17.1%	18.6%	31.6%	9.1%	7.1%	3.9%	0.4%	35.1%***	60.2%
[-10, -5)	392,409	0.7%	9.4%	16.8%	17.5%	29.5%	11.8%	8.8%	5.0%	0.5%	39.7%***	62.4%
[-5, -3)	605,015	0.3%	7.9%	16.1%	16.7%	31.1%	12.4%	10.1%	4.9%	0.5%	42.6%***	60.3%
[-3, -1)	624,552	0.3%	6.0%	13.1%	17.3%	38.1%	12.5%	8.2%	4.2%	0.3%	42.9%***	49.8%
[-1, 1]	1,217,733	0.2%	3.8%	8.7%	13.7%	51.6%	12.1%	6.3%	3.3%	0.3%	47.0%***	53.5%
(1, 3]	476,239	0.3%	5.4%	11.5%	14.3%	35.4%	15.4%	10.7%	6.3%	0.7%	51.0%***	49.0%
(3, 5]	351,814	0.3%	6.2%	13.9%	14.2%	25.6%	15.6%	14.1%	8.9%	1.3%	52.9%***	51.9%
(5, 10]	237,153	0.3%	5.9%	11.9%	13.1%	21.4%	16.0%	14.3%	14.3%	2.8%	58.4%***	53.2%
> 10	29,329	0.5%	5.9%	10.9%	10.9%	15.9%	13.7%	13.4%	18.8%	9.9%	63.9%***	49.4%

Table 8: 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 Trades* is the average number of executed trades per user within the bucket. *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 at the bucket level as $(\overline{\text{PnL}}_{\text{excl.}} - \overline{\text{PnL}}_{\text{incl.}})/|\overline{\text{PnL}}_{\text{incl.}}|$ from the bucket-mean PnL columns rather than by averaging per-user fractions, so the statistic is well-defined even when some users in a bucket have very small absolute PnL. Positive values indicate that the spread reduced the bucket’s PnL; negative values indicate that the bucket is a net spread earner. Users with PnL = 0 are absent from both panels by construction. *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 Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)							
0–20	341,105	30	-0.37	0.36	-0.73***	195.8%	13.1%
20–40	341,105	53	-3.23	0.36	-3.60***	111.1%	11.0%
40–60	341,105	85	-13.58	-5.45	-8.13***	59.8%	10.1%
60–80	341,105	160	-60.87	-50.65	-10.22***	16.8%	14.0%
80–90	170,552	435	-256.41	-235.79	-20.62***	8.0%	17.9%
90–95	85,276	790	-892.34	-859.76	-32.58***	3.7%	28.6%
95–99	68,221	1,468	-2,760.45	-2,690.62	-69.83***	2.5%	35.4%
99–99.9	15,350	4,428	-16,675.11	-16,359.73	-315.38***	1.9%	29.3%
99.9–100	1,705	13,730	-269,401.48	-266,976.23	-2,425.25	0.9%	37.3%
Panel B: Winners (PnL > 0)							
0–20	153,238	29	0.18	0.47	-0.29***	160.3%	24.4%
20–40	153,238	53	2.14	3.12	-0.98***	45.6%	23.1%
40–60	153,237	98	13.27	15.14	-1.87***	14.1%	20.6%
60–80	153,238	192	72.33	75.05	-2.73***	3.8%	24.9%
80–90	76,619	595	344.79	352.91	-8.12***	2.4%	26.3%
90–95	38,309	1,296	1,306.33	1,302.12	4.21	-0.3%	35.1%
95–99	30,648	4,141	5,108.26	5,044.72	63.54***	-1.2%	40.6%
99–99.9	6,895	37,254	38,598.91	37,160.71	1,438.21***	-3.7%	44.8%
99.9–100	766	255,314	701,648.09	678,885.67	22,762.42***	-3.2%	49.3%

Table 9: Simulation Counterfactual Summary

This table reports cross-scenario summary statistics from the simulation benchmark of Section 7. The first row reports the empirical facts from the corresponding panels of Sections 4.1–6.3. The remaining rows are based on the *Simulation baseline* (no information heterogeneity, no overconfidence, homogeneous maker propensity, constant half-spread). The next three rows each include *one* channel on top of this baseline in isolation: *+ Skilled traders* adds dispersion in true signal precision τ_i ; *+ Maker specialization* adds dispersion in maker propensity μ_i ; *+ Overconfidence* adds the perceived-precision wedge $\hat{\tau}_i \neq \tau_i$. The final row, *+ All three*, includes all three channels simultaneously, with all primitives drawn independently across traders. Columns are: the fraction of simulated users with negative cumulative PnL; the share of total winner PnL captured by the top 1% and top 5% of winners; the volume-weighted maker share of the top 1% of earners; the share of total loser PnL borne by the bottom 1% and bottom 5% of losers; the difference in mean per-trade EHR between the top 5% and bottom 5% of users by PnL; and the fraction of users with negative cumulative PnL whose PnL would have been positive had each of their executed shares cleared at the market midpoint instead of at the bid or ask. This last statistic is a same-trade ex-post adjustment computed within each simulation run; it mirrors the empirical spread decomposition in Section 6.2 and does not require a separate counterfactual simulation. The simulation uses $N = 100,000$ traders and $M = 50,000$ markets.

Scenario	Fraction Losers	Top 1% Gains	Top 5% Gains	Maker Sh. Top 1%	Bottom 1% Losses	Bottom 5% Losses	EHR Top–Bot (5% Groups)	Losers Flipped (No Spread)
Empirical Facts	69.0%	76.5%	91.4%	42.0%	68.1%	86.1%	0.280	18.5%
Simulation Baseline	52.0%	30.2%	58.2%	34.1%	32.3%	59.0%	0.006	0.5%
+ Skilled Traders	67.4%	42.7%	70.3%	46.5%	35.4%	61.8%	0.027	0.5%
+ Maker Specialization	61.6%	45.4%	70.6%	70.0%	32.5%	59.7%	0.009	5.3%
+ Overconfidence	55.7%	37.9%	65.0%	33.8%	31.7%	58.9%	0.006	0.4%
+ All Three	78.1%	76.0%	89.4%	66.0%	28.3%	55.1%	0.029	3.4%

Internet Appendix 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

Table of Contents

A	Data Collection	1
	A.1 Reconciling trades from order-filled events.....	2
	A.2 Determining the winning outcome.....	3
	A.3 Computing point-in-time user positions and PnL	5
	A.4 Assigning markets to categories	7
B	Market Horizon	9
	B.1 Distribution of market horizon at inception.....	9
	B.2 Calibration at longer horizons.....	11
C	Wash Trading	13
	C.1 Identification procedure	13
	C.2 Summary statistics.....	14
	C.3 PnL concentration excluding wash traders.....	17
	C.4 Loss-determinants probit excluding wash traders	17
D	Resolved Markets Only	21
E	No-Fee Markets Only	24
F	Long-Horizon PnL Transition Matrix	27
G	Risk-Adjusted PnL Transition Matrices	29
H	PnL Spread Decomposition by Category	34
I	Simulation Benchmark Additional Exhibits	42
J	Heterogeneity of Winners and Losers	47

Internet 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.³⁸

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.2.

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

³⁸Polymarket’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. 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.3. 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.2. 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.³⁹ For the small share of resolved markets whose Polymarket-API resolution price is missing from our local prices stream, we fall back to the token’s last observed trade price as the settlement price; under the one-USDC settlement rule the last trade price is typically already at zero or one for these tokens, so this fallback materially affects only ambiguous edge cases.

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

³⁹The 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.

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

A.4. 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.

B. Market Horizon

B.1. Distribution of market horizon at inception

Table B1: Distribution of Market Horizon at Inception

This table reports the distribution of market horizons at inception, defined as the elapsed time between a market’s open (`market_start_time`) and its resolution (`close_time`). The sample is restricted to markets that have resolved by the end of our sample period. Each row corresponds to one of the seven market categories or to the aggregate (“All”). The reported statistics are the mean, standard deviation, minimum, 5th, 25th, 50th, 75th, and 95th percentiles, and maximum. The unit shown in each column header is selected from the typical magnitude of values in that column: “wk” for weeks, “d” for days, “hr” for hours, and “min” for minutes. The sample period is from November 11, 2022 to March 29, 2026.


Category	Mean (wk)	Std. Dev. (wk)	Min (hr)	5th Pct. (hr)	25th Pct. (d)	Median (d)	75th Pct. (wk)	95th Pct. (wk)	Max (wk)
Sports	1.95	4.44	0.66	17.0	1.20	4.88	1.94	6.35	51.3
Crypto	0.91	2.50	2.26	3.93	0.34	6.62	1.02	1.97	64.8
Politics	4.71	8.55	2.75	12.3	2.70	8.38	4.89	23.8	134
Finance	2.50	4.83	2.18	34.4	6.03	7.07	3.28	6.70	52.8
Tech	4.24	8.06	1.30	37.1	4.81	7.16	4.56	20.1	54.2
Culture	3.48	5.56	2.21	29.0	6.04	10.4	3.35	16.4	52.7
Weather	0.52	1.24	2.65	28.7	2.15	2.70	0.54	0.70	52.2
All	2.09	4.95	0.66	8.11	1.30	5.36	1.70	8.63	134

B.2. Calibration at longer horizons

Section 4 presents calibration at horizons of one and seven days before resolution, using a consistent-sample restriction so that the same market enters every horizon. This appendix relaxes both constraints. Figure B1 extends the horizons to one, two, and three weeks, one month, and three months before close, and uses a per-horizon sample: for each horizon h , the calibration line is computed from all markets that were opened at least h before their resolution and have a reconstructed price at time $\text{close} - h$. The sample therefore grows monotonically as the horizon shrinks. The right panel plots a kernel-density estimate of market lifetimes for the full set of resolved markets, with vertical dashed lines at each horizon used in the left panel; the x -axis is on a log scale. Together the two panels show how much of the sample is available at each horizon and how stable calibration remains as we move away from the resolution date.

Figure B1: Calibration at Longer Horizons and Distribution of Market Lifetimes

Panel A plots the calibration of Polymarket binary market prices at horizons of three months, one month, three weeks, two weeks, and one week before resolution. Market prices are binned into ten equal-width bins; for each bin the x -axis (Expected) shows the bin midpoint and the y -axis (Actual) shows the observed “Yes” frequency, matching the axis convention used in Figure 2. The dashed 45-degree line represents perfect calibration. Shaded bands are 95% bootstrap confidence intervals for the observed frequency (1,000 replications). For each horizon h the sample includes all resolved markets that were opened at least h before close and have a valid reconstructed price at time close $- h$, so the sample grows as the horizon shrinks (it is *not* fixed across horizons). Panel B plots a kernel-density estimate of market lifetimes (close minus open) over all resolved markets, on a log x -axis, with vertical dashed black lines marking six lifetime landmarks (8 hours, 1 day, 7 days, 14 days, 30 days, 182 days). The sample period is from November 11, 2022 to March 29, 2026.



C. Wash Trading

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 when excluding potential wash traders.

C.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:

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

A wallet that routes all of its volume through a single counterparty has $\text{HHI}_i = 1$, while a wallet that splits volume uniformly across K counterparties has $\text{HHI}_i = 1/K$. 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 (Siroly, Ma, Kanoria, and Sethi, 2025).

In what follows we flag a user as a suspected wash trader when $\text{HHI}_i \geq 0.5$. This cutoff is reached when one counterparty accounts for at least roughly 71% of the user’s USDC volume on its own ($\sqrt{0.5} \approx 0.707$), when two counterparties each account for 50%, or by similar concentrated combinations. 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 Siroly, 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 $\sim 25\%$ 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.

C.2. Summary statistics

Table C2 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. Flagging statistics by HHI cutoff (number of wallets, share of users, and share of volume) are reported in Table C3.

Table C2: 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. Two HHI variants are reported: *HHI Count* (shares based on number of trades per counterparty) and *HHI Volume* (shares based on USDC volume; under the matched-notional convention this coincides with token-quantity weighting). For each variant the table reports the number of wallets (N), the cross-wallet mean, and the 25th, 50th, 75th, 90th, 95th, and 99th percentiles. The sample period is from November 11, 2022 to March 29, 2026.

Counterparty HHI	N	Mean	P25	Median	P75	P90	P95	P99
Count-weighted	482,092	0.028	0.010	0.015	0.025	0.048	0.084	0.264
Volume-weighted	482,092	0.073	0.017	0.032	0.071	0.182	0.277	0.628

Table C3 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. Applying the 0.5 cutoff flags 7,241 wallets (1.5% of screened users) and captures 3.1% of 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 $\sim 25\%$ volume share reported by [Siroly, 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 C3: 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 total volume (in millions of USDC) 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). *Volume (M USDC)* counts each trade once whenever at least one side is flagged. The sample period is from November 11, 2022 to March 29, 2026.

HHI Threshold	Users Flagged	% Users	Trades	% Trades	Volume (M)	% Volume
0.9	2,259	0.5%	471,564	0.1%	1,374.2	2.0%
0.7	3,989	0.8%	826,938	0.1%	1,661.4	2.5%
0.5	7,241	1.5%	1,488,118	0.3%	2,083.6	3.1%
0.3	20,804	4.3%	4,339,481	0.7%	3,397.9	5.1%

Table C4 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 [Siroly, Ma, Kanoria, and Sethi \(2025\)](#).

Table C4: Wash-Trading Volume Share by Market Category

This table reports, for each market category, the total volume on the platform (in millions of USDC) 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). Trades with no category label are excluded. The sample period is from November 11, 2022 to March 29, 2026.

Category	Volume (M)	Wash % (HHI $\geq$ 0.9)	Wash % (HHI $\geq$ 0.7)	Wash % (HHI $\geq$ 0.5)	Wash % (HHI $\geq$ 0.3)
Sports	25,477.0	3.4%	3.9%	4.6%	6.8%
Politics	21,114.4	1.2%	1.6%	2.1%	4.1%
Crypto	12,762.0	1.2%	1.6%	2.2%	3.3%
Culture	3,473.9	1.0%	1.7%	2.8%	5.3%
Finance	3,107.5	1.0%	1.4%	1.7%	4.7%
Tech	806.9	2.6%	3.0%	3.3%	4.3%
Weather	393.7	1.7%	2.0%	2.4%	3.5%

C.3. PnL concentration excluding wash traders

Figure C2 replicates the main-text concentration figure (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. The time-series panel and the fraction of users with positive PnL are likewise unchanged. The extreme concentration documented in Section 4.1 is therefore not driven by wash-trading wallets.


C.4. Loss-determinants probit excluding wash traders

Table C5 re-estimates the specifications of Table 4 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 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

Figure C2: 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 ($\text{HHI} \geq 0.5$, volume-weighted; at least 100 settled trades). Panel (a) plots the cumulative concentration as a snapshot at the end of the sample: 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, and the red bar reports the share of total losses borne by users in the corresponding bottom group. Panel (b) plots the same shares over time at a monthly frequency, with one line per percentile bucket (0.1%, 1%, 5%) for gains and losses. Panel (c) plots the fraction of users with positive PnL at each month-end. The sample period is from November 11, 2022 to March 29, 2026.



expect the opposite since the variable proxies for wash-trading, 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 negative profits is more significant than wash trading captures: 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 5, are unaffected.

Table C5: Determinants of Negative Performance Probit Regression: Wash Traders Excluded

This table replicates Table 4 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 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*** (0.001)	−0.017*** (0.001)	−0.007*** (0.001)	−0.003*** (0.001)	0.002 (0.002)	0.007*** (0.002)	−0.129*** (0.006)	−0.128*** (0.006)
Frac Maker Volume		−0.354*** (0.001)	−0.349*** (0.001)	−0.343*** (0.001)	−0.318*** (0.003)	−0.314*** (0.003)	−0.142*** (0.006)	−0.137*** (0.006)
Log N Trades		−0.001*** (0.000)	0.004*** (0.000)	0.018*** (0.000)	0.030*** (0.001)	0.034*** (0.001)	−0.002 (0.002)	−0.005** (0.002)
Log Total Volume		0.020*** (0.000)	0.018*** (0.000)	0.020*** (0.000)	−0.022*** (0.000)	−0.023*** (0.001)	−0.044*** (0.001)	−0.041*** (0.001)
Category HHI			0.080*** (0.001)	−0.057*** (0.002)	−0.019*** (0.003)	−0.080*** (0.004)	−0.136*** (0.007)	−0.146*** (0.010)
Counterparty HHI			−0.002 (0.002)	0.040*** (0.002)	0.354*** (0.010)	0.403*** (0.011)	0.827*** (0.063)	0.812*** (0.063)
Traded Crypto				−0.074*** (0.001)		−0.018*** (0.002)		0.056*** (0.006)
Traded Finance				−0.034*** (0.001)		−0.019*** (0.002)		−0.009* (0.005)
Traded Politics				−0.056*** (0.001)		−0.000 (0.002)		0.041*** (0.005)
Traded Tech				−0.034*** (0.001)		−0.023*** (0.002)		−0.026*** (0.005)
Traded Culture				0.002** (0.001)		−0.005** (0.002)		−0.023*** (0.005)
Traded Weather				−0.036*** (0.001)		−0.024*** (0.002)		−0.008* (0.005)
Observations	2,472,863	2,472,863	2,472,863	2,472,863	470,965	470,965	78,223	78,223
Pseudo R^2	0.000	0.034	0.036	0.043	0.042	0.043	0.055	0.057

D. 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 D6: 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 ($\text{PnL} < 0$) and winners ($\text{PnL} > 0$), 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	614,883	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 (M USDC)	67,135.4	25,477.0	21,114.4	12,762.0	3,473.9	3,107.5	806.9	393.7
Frac. Winners	35.1%	33.1%	39.5%	42.2%	40.7%	47.7%	52.1%	52.5%
Mean PnL by Percentile – Losers ($\text{PnL} < 0$)								
0–20	-6.16	-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	-921	-74	-98	-190	-47	-55	-43	-53
80–90	-3,263	-321	-447	-770	-192	-223	-164	-164
90–95	-8,334	-1,044	-1,469	-2,062	-673	-754	-539	-467
95–99	-30,886	-3,586	-6,227	-6,334	-2,918	-3,035	-1,975	-1,710
99–99.9	-215,599	-22,265	-44,890	-28,933	-21,748	-22,103	-11,759	-8,772
99.9–100	-3,900,155	-353,640	-516,041	-188,674	-278,387	-219,727	-106,754	-59,241
Mean PnL by Percentile – Winners ($\text{PnL} > 0$)								
0–20	1.91	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	548	97	71	49	25	23	13	10
80–90	2,274	449	289	224	94	104	51	37
90–95	6,246	1,440	919	891	294	338	155	117
95–99	24,635	4,787	3,745	4,659	1,416	1,477	787	694
99–99.9	220,121	31,766	27,710	25,642	10,090	9,893	5,606	4,824
99.9–100	5,030,873	573,771	637,746	222,703	177,627	181,434	59,259	55,490

Table D7: Probit Regression: Determinants of Negative Performance — Resolved Markets Only

This table replicates Table 4 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 4. 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*** (0.001)	−0.161*** (0.001)	−0.175*** (0.001)	−0.180*** (0.001)	−0.046*** (0.002)	−0.050*** (0.002)	0.047*** (0.007)	0.049*** (0.007)
Frac Maker Volume		−0.213*** (0.001)	−0.231*** (0.001)	−0.228*** (0.001)	−0.111*** (0.003)	−0.108*** (0.003)	0.079*** (0.006)	0.081*** (0.006)
Log N Trades		0.010*** (0.000)	−0.014*** (0.000)	−0.018*** (0.000)	−0.028*** (0.001)	−0.023*** (0.001)	−0.011*** (0.002)	−0.011*** (0.002)
Log Total Volume		0.021*** (0.000)	0.025*** (0.000)	0.023*** (0.000)	0.014*** (0.000)	0.009*** (0.001)	−0.002 (0.001)	−0.003** (0.001)
Category HHI			−0.089*** (0.001)	0.000 (0.002)	−0.099*** (0.003)	−0.032*** (0.004)	−0.116*** (0.008)	−0.092*** (0.011)
Counterparty HHI			−0.159*** (0.002)	−0.175*** (0.002)	0.019*** (0.007)	0.022*** (0.007)	0.337*** (0.044)	0.350*** (0.044)
Traded Crypto				0.008*** (0.001)		−0.002 (0.002)		0.025*** (0.007)
Traded Finance				0.016*** (0.001)		−0.005*** (0.002)		−0.020*** (0.005)
Traded Politics				0.044*** (0.001)		0.060*** (0.002)		0.020*** (0.005)
Traded Tech				0.018*** (0.001)		0.005*** (0.002)		0.018*** (0.006)
Traded Culture				0.023*** (0.001)		0.015*** (0.002)		0.015*** (0.005)
Traded Weather				0.007*** (0.001)		0.003* (0.002)		−0.019*** (0.005)
Observations	2,480,080	2,480,080	2,480,080	2,480,080	478,096	478,096	78,366	78,366
Pseudo R^2	0.010	0.026	0.032	0.033	0.008	0.009	0.006	0.007

E. 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 E8: 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	614,883	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 (M USDC)	67,135.4	25,477.0	21,114.4	12,762.0	3,473.9	3,107.5	806.9	393.7
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	-6.67	-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	-1,047	-55	-53	-161	-30	-33	-31	-48
80–90	-3,596	-249	-275	-748	-132	-147	-118	-153
90–95	-9,211	-858	-1,029	-2,161	-480	-531	-386	-427
95–99	-33,751	-3,135	-4,887	-7,131	-2,351	-2,472	-1,665	-1,613
99–99.9	-224,945	-20,391	-37,095	-44,966	-19,200	-18,888	-10,682	-8,117
99.9–100	-3,144,881	-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	539	64	27	95	15	12	9.09	9.50
80–90	2,217	310	121	599	55	51	33	34
90–95	5,975	1,158	446	2,362	163	163	98	98
95–99	22,751	4,242	2,458	9,856	882	888	501	581
99–99.9	246,755	32,954	23,712	55,846	8,434	7,537	5,046	4,051
99.9–100	7,510,162	633,831	629,239	499,640	164,905	171,593	56,974	50,539

Table E9: Probit Regression: Determinants of Negative Performance — No-Fee Markets Only

This table replicates Table 4 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 4. 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	−0.186*** (0.001)	−0.206*** (0.001)	−0.228*** (0.001)	−0.239*** (0.001)	−0.156*** (0.002)	−0.179*** (0.002)	−0.030*** (0.007)	−0.043*** (0.007)
Frac Maker Volume		−0.314*** (0.001)	−0.337*** (0.001)	−0.318*** (0.001)	−0.304*** (0.003)	−0.294*** (0.003)	−0.156*** (0.006)	−0.154*** (0.006)
Log N Trades		0.013*** (0.000)	−0.014*** (0.000)	−0.006*** (0.000)	−0.028*** (0.001)	−0.006*** (0.001)	−0.028*** (0.002)	−0.012*** (0.002)
Log Total Volume		0.029*** (0.000)	0.034*** (0.000)	0.029*** (0.000)	0.031*** (0.000)	0.013*** (0.001)	0.013*** (0.001)	−0.004*** (0.001)
Category HHI			−0.154*** (0.001)	−0.090*** (0.002)	−0.265*** (0.003)	−0.143*** (0.004)	−0.347*** (0.008)	−0.169*** (0.011)
Counterparty HHI			−0.137*** (0.002)	−0.130*** (0.002)	−0.024*** (0.007)	0.021*** (0.007)	0.322*** (0.044)	0.359*** (0.046)
Traded Crypto				−0.051*** (0.001)		−0.095*** (0.002)		−0.098*** (0.007)
Traded Finance				−0.009*** (0.001)		0.010*** (0.002)		0.028*** (0.005)
Traded Politics				0.077*** (0.001)		0.158*** (0.002)		0.142*** (0.005)
Traded Tech				−0.012*** (0.001)		−0.019*** (0.002)		−0.033*** (0.006)
Traded Culture				0.042*** (0.001)		0.026*** (0.002)		0.046*** (0.005)
Traded Weather				−0.011*** (0.001)		0.003 (0.002)		0.013*** (0.005)
Observations	2,480,080	2,480,080	2,480,080	2,480,080	478,096	478,096	78,366	78,366
Pseudo R^2	0.016	0.047	0.055	0.061	0.042	0.056	0.035	0.049

F. Long-Horizon PnL Transition Matrix

This appendix extends the monthly PnL transition matrix in Table 6 to longer horizons. The construction is identical to the main-text table: users who traded are sorted into eleven monthly-PnL-change groups; the eleven bucket columns report transition frequencies *conditional on trading* (blue-shaded); the “> 0” column reports the conditional probability of a strictly positive target-month PnL change with significance stars from a two-sided exact-binomial test against the null hypothesis that the probability of profit equals 0.5 (cells shaded only when significant — green above 0.5, red below 0.5); and the rightmost (blue-shaded) column reports the unconditional “No Trade” rate.

Table F10: Transition Matrix of PnL Groups — Long Horizons

This table reports transition probabilities across monthly-performance groups over 6-month and 12-month horizons. In each month, users who traded are sorted into eleven groups based on the change in their mark-to-market profit and loss (PnL) during that month, using the breakpoints $-100k$, $-10k$, $-1k$, -100 , -10 , 10 , 100 , $1k$, $10k$, and $100k$ (in dollars). The eleven bucket columns report transition frequencies *conditional on trading* (each row's bucket-to-bucket cells sum to 100% and are shaded in blue). The “ > 0 ” column reports the conditional probability that the trader's target-month PnL change is strictly positive. Significance stars come from a two-sided exact-binomial test against the null hypothesis that the probability of profit equals 0.5 (i.e., that profits and losses are equally likely): * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Cells are shaded only when the test rejects the null at the 10% level: green for probabilities above 0.5, red for probabilities below 0.5. The “No Trade” column reports the unconditional share of users in each starting group with no positions in the target month, shaded in blue. The sample period is from November 11, 2022 to March 29, 2026.

<i>Panel A: 6-Month Horizon</i>														
Starting Group	<i>N</i>	$< -100k$	$[-100k, -10k)$	$[-10k, -1k)$	$[-1k, -100)$	$[-100, -10)$	$[-10, 10]$	$(10, 100]$	$(100, 1k]$	$(1k, 10k]$	$(10k, 100k]$	$> 100k$	> 0	No Trade
$< -100k$	492	11.6%	21.9%	9.0%	7.7%	0.6%	1.3%	1.9%	3.9%	13.5%	21.9%	6.5%	49.0%	68.5%
$[-100k, -10k)$	4,890	2.0%	14.1%	21.3%	10.9%	2.9%	5.4%	2.8%	8.3%	16.1%	13.8%	2.4%	46.7%***	61.6%
$[-10k, -1k)$	25,921	0.3%	3.9%	19.7%	18.3%	7.5%	12.4%	5.1%	13.3%	14.8%	4.4%	0.3%	44.0%***	62.8%
$[-1k, -100)$	78,852	0.0%	0.8%	6.9%	20.0%	16.8%	23.9%	12.0%	12.8%	5.8%	0.9%	0.0%	44.9%***	67.4%
$[-100, -10)$	206,858	0.0%	0.1%	1.5%	6.9%	17.2%	56.3%	11.7%	4.9%	1.2%	0.1%	0.0%	51.2%***	69.2%
$[-10, 10]$	2,097,292	0.0%	0.0%	1.2%	1.8%	5.2%	85.9%	4.1%	1.1%	0.4%	0.1%	0.0%	60.7%***	63.5%
$(10, 100]$	153,864	0.0%	0.2%	1.8%	6.4%	14.7%	54.4%	14.9%	5.9%	1.6%	0.2%	0.0%	56.6%***	64.1%
$(100, 1k]$	56,925	0.0%	0.7%	6.8%	16.2%	13.7%	23.3%	12.1%	17.1%	8.8%	1.2%	0.0%	53.4%***	58.7%
$(1k, 10k]$	21,497	0.2%	3.8%	16.2%	14.5%	5.5%	9.4%	4.9%	14.2%	23.4%	7.5%	0.4%	54.8%***	47.8%
$(10k, 100k]$	4,878	2.0%	11.6%	14.2%	7.8%	1.9%	4.5%	1.6%	6.2%	19.0%	26.9%	4.3%	61.2%***	42.3%
$> 100k$	526	5.3%	14.5%	8.8%	3.1%	0.9%	0.9%	0.4%	1.3%	15.9%	25.1%	23.8%	67.0%***	56.8%
<i>Panel B: 12-Month Horizon</i>														
Starting Group	<i>N</i>	$< -100k$	$[-100k, -10k)$	$[-10k, -1k)$	$[-1k, -100)$	$[-100, -10)$	$[-10, 10]$	$(10, 100]$	$(100, 1k]$	$(1k, 10k]$	$(10k, 100k]$	$> 100k$	> 0	No Trade
$< -100k$	293	7.0%	18.3%	11.3%	1.4%	2.8%	5.6%	2.8%	5.6%	16.9%	26.8%	1.4%	56.3%	75.8%
$[-100k, -10k)$	2,797	1.4%	14.4%	18.6%	10.7%	3.7%	3.8%	3.9%	10.5%	17.3%	13.5%	2.3%	50.0%	71.7%
$[-10k, -1k)$	13,712	0.5%	4.7%	17.0%	19.8%	10.1%	9.6%	6.1%	13.3%	13.3%	5.1%	0.5%	43.1%***	71.4%
$[-1k, -100)$	42,078	0.1%	0.8%	7.2%	19.0%	16.6%	24.8%	11.9%	12.0%	6.2%	1.4%	0.1%	46.1%***	73.3%
$[-100, -10)$	122,434	0.0%	0.2%	3.1%	8.8%	14.8%	57.4%	9.6%	4.4%	1.5%	0.3%	0.0%	46.2%***	68.1%
$[-10, 10]$	1,024,314	0.0%	0.1%	1.2%	3.3%	7.0%	81.5%	4.5%	1.9%	0.5%	0.1%	0.0%	58.0%***	69.3%
$(10, 100]$	91,686	0.0%	0.1%	3.3%	8.2%	13.7%	55.0%	11.8%	6.0%	1.5%	0.4%	0.0%	50.4%	66.0%
$(100, 1k]$	32,158	0.0%	0.9%	7.1%	15.5%	13.8%	23.0%	12.7%	15.9%	9.1%	1.8%	0.1%	54.3%***	67.5%
$(1k, 10k]$	11,110	0.3%	4.2%	14.4%	16.4%	6.8%	8.1%	5.4%	14.3%	20.4%	9.1%	0.8%	53.7%***	59.2%
$(10k, 100k]$	2,604	1.6%	11.4%	15.3%	8.8%	1.7%	2.1%	1.4%	8.5%	19.4%	24.7%	5.1%	60.5%***	54.2%
$> 100k$	302	8.0%	14.3%	8.9%	5.4%	0.9%	3.6%	1.8%	6.2%	12.5%	25.0%	13.4%	61.6%**	62.9%

G. Risk-Adjusted PnL Transition Matrices

This appendix complements the dollar-PnL transition matrix in Table 6 and the Sharpe-analog risk-adjusted transition matrix in Table 7 with two additional sets of risk-adjusted transition matrices. The construction otherwise mirrors Table 7: in each month, users with sufficient activity are sorted into nine groups, and we report (i) the percentage of users in each starting group whose ratio falls into each group at the specified horizon, conditional on a valid target-month observation, (ii) the “> 0” column, which gives the conditional probability of a strictly positive target-month ratio with significance stars from a two-sided exact-binomial test against $H_0: p = 0.5$, and (iii) the unconditional “No Trade” rate.

For each user-month we compute the daily mark-to-market PnL change on every active day, where a day is active iff the user traded that day or held a non-zero non-USDC position at the end of the prior day. We drop user-months with fewer than 10 active days. For the remaining user-months we compute the mean and standard deviation of the daily PnL changes (using $\text{ddof} = 1$), as well as a downside semi-deviation defined as the root-mean-square of the negative daily changes (treating non-negative observations as zero). We annualize each ratio by $\sqrt{365}$, since Polymarket trades seven days a week. We then bucket users into nine groups using cutoffs at ± 10 , ± 5 , ± 3 , and ± 1 , leaving a single merged middle bucket $[-1, 1]$ that captures all small-magnitude ratios. User-months whose relevant standard deviation is at or below \$0.10 are dropped, since the resulting ratio is dominated by floating-point noise on near-flat panels. The merged middle bucket spans both signs of the ratio, so we compute the “> 0” statistic directly from the raw target-month ratio rather than from the bucket index.

We report two ratios: the PnL-to-volatility ratio (Sharpe analog), $\sqrt{365} \cdot \overline{\Delta \text{PnL}_d} / \sigma(\Delta \text{PnL}_d)$, and the PnL-to-downside-volatility ratio (Sortino analog), which replaces the denominator with the downside semi-deviation. As with the user-level Sharpe and Sortino analogs reported in Table 5,

these are constructed on dollar PnL rather than percentage returns and are not strictly Sharpe or Sortino ratios. Table G11 extends the main-text Table 7 (Sharpe analog) to 6- and 12-month horizons. Tables G12 and G13 report the Sortino-analog counterparts at 1-/3-month and 6-/12-month horizons, respectively.

Table G11: Risk-Adjusted PnL Transition Matrix — Sharpe Analog, Long Horizons

This table extends Table 7 to 6-month and 12-month horizons. Construction is otherwise identical.

<i>Panel A: 6-Month Horizon</i>												
Starting Group	<i>N</i>	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	21,321	1.2%	12.6%	15.5%	17.9%	29.3%	11.5%	7.3%	4.3%	0.4%	37.1%***	77.5%
[-10, -5)	275,640	0.6%	9.2%	16.1%	16.6%	29.8%	12.3%	9.4%	5.3%	0.5%	42.4%***	76.1%
[-5, -3)	445,539	0.5%	7.8%	15.0%	16.0%	31.3%	12.8%	10.3%	5.8%	0.6%	44.9%***	75.0%
[-3, -1)	457,811	0.4%	6.8%	13.6%	16.1%	36.0%	12.5%	8.7%	5.6%	0.5%	44.5%***	69.6%
[-1, 1]	863,127	0.2%	4.9%	11.1%	15.1%	42.8%	12.6%	8.1%	4.6%	0.5%	46.2%***	73.4%
(1, 3]	341,030	0.3%	6.2%	11.9%	15.2%	32.6%	15.3%	11.1%	6.7%	0.8%	50.2%	65.6%
(3, 5]	253,066	0.3%	6.6%	12.7%	14.1%	28.0%	16.0%	12.6%	8.5%	1.2%	52.5%***	65.6%
(5, 10]	163,246	0.3%	6.1%	11.4%	12.4%	23.8%	15.5%	13.6%	13.9%	2.8%	58.5%***	65.6%
> 10	12,627	0.3%	4.4%	8.5%	10.2%	19.1%	15.8%	13.5%	20.9%	7.2%	68.3%***	59.7%
<i>Panel B: 12-Month Horizon</i>												
Starting Group	<i>N</i>	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	14,543	2.1%	13.6%	16.9%	15.5%	24.6%	12.7%	8.0%	5.6%	1.0%	37.8%***	89.8%
[-10, -5)	192,766	0.8%	9.9%	17.1%	15.5%	25.2%	12.4%	11.0%	7.0%	1.1%	42.2%***	84.9%
[-5, -3)	299,964	0.6%	8.5%	18.1%	15.3%	23.2%	13.2%	12.4%	7.4%	1.3%	44.9%***	83.5%
[-3, -1)	278,666	0.7%	8.3%	16.4%	16.0%	25.1%	13.5%	11.2%	7.8%	1.1%	45.3%***	81.0%
[-1, 1]	441,280	0.5%	7.4%	15.2%	15.2%	27.6%	13.8%	11.8%	7.6%	1.0%	47.4%***	83.4%
(1, 3]	182,683	0.5%	7.5%	14.0%	14.8%	23.6%	16.0%	13.2%	8.9%	1.5%	51.8%***	76.8%
(3, 5]	151,535	0.5%	7.3%	15.2%	14.8%	21.4%	14.7%	14.2%	10.2%	1.7%	51.7%***	75.8%
(5, 10]	98,952	0.5%	6.5%	13.2%	13.5%	20.5%	15.4%	14.9%	12.9%	2.6%	56.4%***	75.6%
> 10	6,479	0.3%	4.6%	9.7%	13.0%	20.9%	15.9%	13.6%	16.6%	5.2%	61.9%***	72.3%

Table G12: Risk-Adjusted PnL Transition Matrix — Sortino Analog

This table reports transition probabilities across monthly risk-adjusted-performance groups over 1-month and 3-month horizons, using the downside-volatility (Sortino-analog) ratio. In each month, users with at least 10 active days are sorted into nine groups based on the annualized $\sqrt{365} \cdot \Delta \text{PnL}_d / \sigma_{\downarrow}(\Delta \text{PnL}_d)$ ratio, where $\sigma_{\downarrow}$ is the root-mean-square of the negative daily PnL changes (treating non-negative observations as zero), with cutoffs at ± 10 , ± 5 , ± 3 , and ± 1 . Bucket cells, the “> 0” column, and the No Trade column have the same interpretation as in Table 7. Long-horizon results (6-month and 12-month) are reported in Table G13. The sample period is from November 11, 2022 to March 29, 2026.

<i>Panel A: 1-Month Horizon</i>												
Starting Group	<i>N</i>	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	34,123	6.0%	27.0%	18.6%	16.0%	13.7%	6.5%	2.9%	3.6%	5.8%	25.2%***	45.8%
[-10, -5)	602,221	0.7%	17.1%	21.5%	18.8%	18.4%	8.2%	4.0%	5.0%	6.4%	31.9%***	46.0%
[-5, -3)	823,027	0.3%	12.2%	21.2%	18.2%	21.9%	9.5%	4.5%	5.4%	6.9%	36.5%***	44.8%
[-3, -1)	709,970	0.2%	8.5%	15.0%	16.8%	29.4%	13.5%	4.9%	5.5%	6.3%	45.5%***	28.8%
[-1, 1]	1,009,089	0.1%	5.1%	9.5%	16.5%	46.9%	11.0%	3.4%	3.6%	4.0%	44.0%***	28.8%
(1, 3]	396,864	0.1%	7.8%	13.1%	21.6%	29.6%	9.3%	5.1%	6.4%	6.9%	39.3%***	27.5%
(3, 5]	192,530	0.2%	8.8%	16.3%	18.3%	18.9%	10.8%	6.9%	9.0%	10.9%	46.3%***	27.6%
(5, 10]	248,636	0.2%	9.4%	15.9%	15.2%	17.7%	11.1%	7.0%	10.1%	13.3%	50.2%*	29.4%
> 10	351,242	0.3%	9.7%	17.0%	14.2%	15.3%	10.0%	6.6%	10.0%	16.9%	50.8%***	33.5%
<i>Panel B: 3-Month Horizon</i>												
Starting Group	<i>N</i>	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	23,415	1.6%	13.2%	20.3%	19.1%	26.1%	7.3%	3.1%	3.8%	5.7%	31.5%***	61.1%
[-10, -5)	469,813	0.5%	13.5%	19.8%	17.9%	22.4%	8.9%	4.5%	5.6%	7.0%	35.7%***	64.2%
[-5, -3)	667,953	0.2%	10.3%	19.9%	17.8%	24.6%	9.8%	4.7%	5.3%	7.3%	38.5%***	63.5%
[-3, -1)	588,796	0.2%	7.9%	15.8%	18.3%	31.7%	10.6%	4.3%	4.9%	6.3%	40.4%***	54.7%
[-1, 1]	851,924	0.1%	5.1%	10.7%	16.6%	43.1%	11.9%	3.9%	3.9%	4.8%	44.9%***	57.8%
(1, 3]	329,465	0.1%	6.8%	12.7%	16.0%	33.3%	11.9%	5.4%	6.4%	7.3%	47.4%***	53.5%
(3, 5]	158,117	0.2%	7.9%	14.9%	15.6%	24.3%	12.2%	6.6%	8.7%	9.6%	48.6%***	50.5%
(5, 10]	201,140	0.2%	8.0%	15.8%	15.5%	20.6%	11.7%	7.2%	9.6%	11.3%	49.9%	49.7%
> 10	273,364	0.2%	8.6%	16.0%	14.8%	18.8%	10.5%	6.6%	9.6%	14.9%	50.6%***	52.4%

Table G13: Risk-Adjusted PnL Transition Matrix — Sortino Analog, Long Horizons

This table extends Table G12 to 6-month and 12-month horizons. Construction is otherwise identical.

<i>Panel A: 6-Month Horizon</i>												
Starting Group	<i>N</i>	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	13,777	1.4%	16.3%	18.8%	17.8%	22.5%	9.3%	3.5%	4.3%	6.1%	32.9%***	78.5%
[-10, -5)	329,064	0.4%	13.7%	19.0%	17.0%	23.0%	9.4%	4.8%	5.8%	6.8%	38.1%***	77.6%
[-5, -3)	494,001	0.2%	10.6%	18.3%	17.3%	24.9%	10.0%	4.9%	5.9%	7.9%	40.5%***	77.7%
[-3, -1)	429,983	0.2%	9.7%	16.8%	17.0%	28.5%	10.4%	4.6%	5.7%	7.0%	41.1%***	73.9%
[-1, 1]	598,100	0.1%	6.9%	14.1%	18.0%	34.2%	10.7%	4.5%	5.3%	6.1%	42.8%***	76.4%
(1, 3]	234,489	0.2%	7.7%	14.2%	17.1%	28.2%	11.4%	5.9%	7.2%	8.1%	46.1%***	71.3%
(3, 5]	111,757	0.2%	9.4%	14.5%	16.0%	23.7%	11.7%	6.7%	8.5%	9.5%	47.8%***	66.3%
(5, 10]	142,563	0.2%	9.6%	14.4%	15.3%	22.3%	11.8%	6.7%	8.6%	11.2%	49.1%***	65.6%
> 10	187,995	0.2%	8.9%	15.4%	14.7%	20.1%	11.1%	6.4%	9.4%	13.7%	50.6%***	66.2%
<i>Panel B: 12-Month Horizon</i>												
Starting Group	<i>N</i>	< -10	[-10, -5)	[-5, -3)	[-3, -1)	[-1, 1]	(1, 3]	(3, 5]	(5, 10]	> 10	> 0	No Trade
< -10	9,772	2.2%	18.9%	19.1%	15.9%	19.1%	9.1%	4.5%	4.6%	6.6%	31.9%***	91.8%
[-10, -5)	226,350	0.6%	13.8%	20.8%	18.6%	17.2%	8.7%	5.0%	6.8%	8.4%	36.7%***	86.1%
[-5, -3)	331,097	0.4%	11.5%	22.3%	17.4%	17.9%	9.0%	5.3%	7.0%	9.2%	38.7%***	85.4%
[-3, -1)	255,740	0.4%	11.9%	20.5%	16.9%	18.4%	9.7%	5.6%	7.1%	9.5%	40.3%***	83.8%
[-1, 1]	307,441	0.3%	10.3%	19.1%	15.9%	22.6%	9.8%	5.4%	7.1%	9.4%	42.2%***	85.4%
(1, 3]	121,803	0.4%	10.5%	18.1%	15.5%	18.6%	10.9%	6.4%	9.0%	10.6%	46.0%***	81.3%
(3, 5]	62,230	0.4%	10.1%	16.4%	15.3%	18.0%	11.5%	6.9%	9.5%	12.1%	48.6%***	76.9%
(5, 10]	85,942	0.4%	10.0%	16.7%	15.6%	17.1%	10.6%	6.5%	10.6%	12.6%	48.8%***	76.2%
> 10	118,806	0.4%	10.4%	16.7%	14.3%	16.7%	10.0%	6.6%	10.2%	14.8%	49.7%	75.8%

H. PnL Spread Decomposition by Category

The following tables replicate Table 8 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.

Table H14: 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 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)							
0–20	211,417	57	-0.22	-0.00	-0.21***	97.9%	14.0%
20–40	211,417	111	-1.88	-0.64	-1.24***	65.9%	10.3%
40–60	211,417	203	-10.47	-7.88	-2.59***	24.8%	13.3%
60–80	211,417	288	-55.70	-51.92	-3.78***	6.8%	16.3%
80–90	105,709	501	-250.75	-243.25	-7.50***	3.0%	16.7%
90–95	52,854	763	-862.53	-847.65	-14.87***	1.7%	25.0%
95–99	42,284	1,186	-3,134.19	-3,090.97	-43.22***	1.4%	28.5%
99–99.9	9,513	3,357	-20,143.64	-19,917.74	-225.90***	1.1%	23.1%
99.9–100	1,057	21,808	-327,713.83	-323,742.90	-3,970.93	1.2%	28.8%
Panel B: Winners (PnL > 0)							
0–20	99,122	128	0.08	0.20	-0.12***	147.6%	22.3%
20–40	99,122	200	2.30	2.77	-0.47***	20.3%	21.8%
40–60	99,122	347	13.04	14.16	-1.12***	8.6%	21.6%
60–80	99,122	352	61.17	64.07	-2.90***	4.7%	23.2%
80–90	49,561	594	293.38	298.27	-4.89***	1.7%	27.4%
90–95	24,780	981	1,108.16	1,117.34	-9.17***	0.8%	35.2%
95–99	19,824	2,267	4,045.72	4,035.42	10.30	-0.3%	44.8%
99–99.9	4,461	14,159	30,626.17	29,881.76	744.41***	-2.4%	44.2%
99.9–100	495	65,714	595,602.95	577,803.98	17,798.98***	-3.0%	52.8%

Table H15: 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 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)							
0–20	154,359	112	-1.07	-0.65	-0.42***	39.1%	15.8%
20–40	154,358	156	-10.19	-9.34	-0.85***	8.3%	15.3%
40–60	154,358	243	-41.29	-39.63	-1.66***	4.0%	16.5%
60–80	154,358	497	-183.41	-178.47	-4.93***	2.7%	18.7%
80–90	77,179	1,047	-756.84	-741.82	-15.02***	2.0%	19.8%
90–95	38,590	1,522	-2,029.71	-2,008.11	-21.61***	1.1%	30.6%
95–99	30,872	2,623	-6,174.81	-6,116.71	-58.10***	0.9%	23.1%
99–99.9	6,946	6,967	-27,862.56	-27,579.61	-282.95***	1.0%	17.9%
99.9–100	771	17,906	-177,033.74	-175,888.43	-1,145.31***	0.6%	17.7%
Panel B: Winners (PnL > 0)							
0–20	116,813	66	0.10	0.14	-0.04***	42.2%	18.9%
20–40	116,812	92	1.15	1.37	-0.22***	18.7%	18.5%
40–60	116,813	123	7.08	7.42	-0.33***	4.7%	21.3%
60–80	116,812	234	34.02	34.93	-0.91***	2.7%	26.6%
80–90	58,406	631	165.75	167.36	-1.61***	1.0%	35.4%
90–95	29,203	1,884	685.60	685.31	0.29	-0.0%	40.8%
95–99	23,363	5,534	4,146.07	4,097.22	48.85***	-1.2%	48.7%
99–99.9	5,256	22,317	23,930.61	23,500.35	430.26***	-1.8%	58.6%
99.9–100	584	415,149	217,583.35	209,525.89	8,057.46***	-3.7%	56.9%

Table H16: 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 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)							
0–20	196,042	145	-0.22	0.15	-0.37***	169.2%	21.3%
20–40	196,041	166	-2.09	0.17	-2.27***	108.2%	12.0%
40–60	196,042	225	-9.32	0.39	-9.70***	104.1%	11.8%
60–80	196,041	305	-52.62	-43.33	-9.28***	17.6%	13.1%
80–90	98,021	447	-275.15	-263.39	-11.76***	4.3%	12.2%
90–95	49,010	633	-1,029.40	-1,012.26	-17.15***	1.7%	13.1%
95–99	39,208	1,220	-4,885.85	-4,835.35	-50.50***	1.0%	13.5%
99–99.9	8,822	3,483	-37,090.00	-37,008.49	-81.51	0.2%	15.0%
99.9–100	980	13,339	-441,610.79	-442,994.79	1,384.00**	-0.3%	23.5%
Panel B: Winners (PnL > 0)							
0–20	111,355	103	0.07	0.15	-0.08***	121.6%	27.0%
20–40	111,354	136	0.88	1.23	-0.35***	39.4%	22.3%
40–60	111,355	436	4.94	4.90	0.03	-0.7%	22.5%
60–80	111,354	246	26.94	25.30	1.64***	-6.1%	24.1%
80–90	55,677	456	120.68	117.97	2.72***	-2.3%	21.2%
90–95	27,839	866	444.49	438.66	5.84***	-1.3%	22.3%
95–99	22,271	2,408	2,452.33	2,405.73	46.60***	-1.9%	27.1%
99–99.9	5,011	10,188	23,752.33	23,407.33	345.00***	-1.5%	36.3%
99.9–100	556	28,699	629,959.47	622,123.32	7,836.15***	-1.2%	46.1%

Table H17: 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: $\text{PnL} < 0$) and winners (Panel B: $\text{PnL} > 0$). See Table 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers ($\text{PnL} < 0$)							
0–20	83,820	143	-0.07	0.01	-0.09***	118.9%	21.5%
20–40	83,819	226	-0.85	-0.34	-0.51***	59.9%	14.3%
40–60	83,819	321	-5.26	-4.21	-1.05***	19.9%	15.8%
60–80	83,819	579	-29.74	-28.57	-1.17***	3.9%	19.7%
80–90	41,910	804	-132.13	-130.05	-2.08***	1.6%	18.1%
90–95	20,955	1,292	-479.46	-475.83	-3.63**	0.8%	17.2%
95–99	16,764	2,159	-2,349.76	-2,339.26	-10.50***	0.4%	17.3%
99–99.9	3,771	4,566	-19,200.30	-19,188.44	-11.87	0.1%	16.4%
99.9–100	419	16,355	-263,776.71	-263,379.49	-397.22***	0.2%	13.7%
Panel B: Winners ($\text{PnL} > 0$)							
0–20	56,615	112	0.03	0.07	-0.04***	142.2%	24.1%
20–40	56,614	227	0.42	0.48	-0.06***	14.5%	25.4%
40–60	56,614	280	2.97	3.06	-0.09***	3.1%	24.5%
60–80	56,614	349	14.96	15.27	-0.31***	2.1%	27.0%
80–90	28,307	697	54.76	55.24	-0.48***	0.9%	27.2%
90–95	14,154	1,506	162.24	159.80	2.44***	-1.5%	30.7%
95–99	11,323	5,257	880.81	859.72	21.09***	-2.4%	35.7%
99–99.9	2,547	19,087	8,420.73	8,244.98	175.75***	-2.1%	43.8%
99.9–100	283	32,697	164,904.64	164,522.62	382.02	-0.2%	56.5%

Table H18: 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: $\text{PnL} < 0$) and winners (Panel B: $\text{PnL} > 0$). See Table 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers ($\text{PnL} < 0$)							
0–20	72,454	212	-0.10	-0.03	-0.07***	65.6%	18.1%
20–40	72,453	267	-1.06	-0.59	-0.47***	44.1%	14.5%
40–60	72,454	398	-6.61	-5.74	-0.88***	13.3%	16.9%
60–80	72,453	525	-33.02	-31.09	-1.92***	5.8%	18.2%
80–90	36,227	862	-146.26	-143.62	-2.64***	1.8%	18.6%
90–95	18,113	1,457	-529.71	-526.92	-2.79***	0.5%	19.1%
95–99	14,491	2,025	-2,467.29	-2,456.23	-11.06***	0.4%	19.8%
99–99.9	3,260	6,687	-18,861.03	-18,845.83	-15.20	0.1%	14.6%
99.9–100	362	9,763	-199,153.69	-199,183.68	29.99	-0.0%	15.6%
Panel B: Winners ($\text{PnL} > 0$)							
0–20	67,086	97	0.02	0.07	-0.05**	198.2%	25.2%
20–40	67,086	154	0.24	0.26	-0.02***	9.2%	22.8%
40–60	67,085	266	1.89	1.98	-0.09***	4.5%	22.1%
60–80	67,086	407	11.52	11.66	-0.15***	1.3%	27.1%
80–90	33,543	812	50.87	51.11	-0.24	0.5%	28.5%
90–95	16,771	1,671	162.53	161.89	0.64	-0.4%	32.9%
95–99	13,417	3,355	881.57	871.08	10.49***	-1.2%	36.3%
99–99.9	3,019	13,355	7,512.26	7,396.19	116.07***	-1.5%	44.2%
99.9–100	335	43,211	171,619.83	171,276.06	343.77	-0.2%	55.1%

Table H19: 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 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers (PnL < 0)							
0–20	42,725	221	-0.08	-0.04	-0.04**	52.3%	21.9%
20–40	42,725	379	-1.07	-0.87	-0.21***	19.4%	19.1%
40–60	42,724	559	-6.57	-6.32	-0.25***	3.8%	19.1%
60–80	42,725	616	-30.92	-30.40	-0.52***	1.7%	21.8%
80–90	21,362	1,277	-118.29	-117.06	-1.24***	1.0%	19.7%
90–95	10,681	1,283	-385.17	-381.45	-3.72***	1.0%	19.7%
95–99	8,545	2,885	-1,663.14	-1,653.78	-9.37***	0.6%	21.6%
99–99.9	1,923	4,124	-10,662.74	-10,648.21	-14.53	0.1%	20.8%
99.9–100	213	7,813	-98,723.51	-98,716.73	-6.79	0.0%	18.0%
Panel B: Winners (PnL > 0)							
0–20	47,256	137	0.02	0.02	0.00	-5.5%	20.9%
20–40	47,256	158	0.25	0.28	-0.04***	14.7%	22.1%
40–60	47,255	327	1.52	1.55	-0.03**	1.9%	21.3%
60–80	47,256	450	9.09	9.10	-0.01	0.1%	24.0%
80–90	23,628	752	33.47	33.61	-0.15*	0.4%	29.2%
90–95	11,814	1,302	97.79	97.88	-0.09	0.1%	30.1%
95–99	9,451	4,984	498.98	492.38	6.61***	-1.3%	34.5%
99–99.9	2,126	21,622	5,020.76	4,985.25	35.52**	-0.7%	39.6%
99.9–100	236	31,193	56,911.03	56,532.54	378.49***	-0.7%	55.2%

Table H20: 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: $\text{PnL} < 0$) and winners (Panel B: $\text{PnL} > 0$). See Table 8 for variable definitions. The sample period is from November 11, 2022 to March 29, 2026.

Percentile	N Users	Mean Trades	Mean PnL (incl.)	Mean PnL (excl.)	Difference	Frac. Spread	Maker %
Panel A: Losers ($\text{PnL} < 0$)							
0–20	17,869	511	-0.39	-0.35	-0.04	9.8%	21.8%
20–40	17,869	739	-3.46	-3.16	-0.31***	8.8%	17.8%
40–60	17,868	907	-13.81	-13.56	-0.25***	1.8%	20.1%
60–80	17,869	1,243	-47.65	-46.87	-0.78***	1.6%	21.0%
80–90	8,934	2,234	-152.55	-150.99	-1.57***	1.0%	22.0%
90–95	4,467	4,559	-426.25	-422.41	-3.84***	0.9%	22.6%
95–99	3,574	3,646	-1,612.57	-1,608.05	-4.52*	0.3%	19.2%
99–99.9	804	8,308	-8,116.88	-8,097.49	-19.38**	0.2%	18.9%
99.9–100	89	13,427	-52,959.27	-52,941.04	-18.23	0.0%	19.9%
Panel B: Winners ($\text{PnL} > 0$)							
0–20	20,991	164	0.01	-0.01	0.02	-182.2%	20.3%
20–40	20,990	240	0.16	0.18	-0.02***	9.6%	24.2%
40–60	20,990	389	1.29	1.36	-0.08***	5.9%	20.1%
60–80	20,990	741	9.46	9.57	-0.11***	1.1%	31.4%
80–90	10,495	1,783	34.02	34.26	-0.23***	0.7%	32.0%
90–95	5,248	3,291	97.97	98.19	-0.22	0.2%	38.8%
95–99	4,198	9,031	581.23	575.09	6.13***	-1.1%	43.2%
99–99.9	945	26,985	4,049.56	4,000.05	49.52***	-1.2%	42.3%
99.9–100	104	47,083	50,547.73	50,308.76	238.96	-0.5%	58.0%

I. Simulation Benchmark Additional Exhibits

This appendix collects the additional exhibits produced by the simulation benchmark of Section 7: the model parameter values, descriptive statistics for the trader cross-section under each scenario, the simulated price calibration, the simulated volume-weighted EHR gradient by PnL percentile, and the simulated maker share by PnL percentile.

Table I21: Simulation Benchmark Parameters and Descriptive Statistics

Panel A reports the hand-tuned model parameters used for the simulation benchmark in Section 7. The same values apply to all specifications except for the scenario-specific overrides described in Section A.8. Panel B reports cross-sectional descriptive statistics of the simulated trader primitives, by specification: the cross-trader mean and standard deviation of $\log W_i$, $\log \tau_i$, logit μ_i , and the overconfidence wedge b_i . Each primitive is drawn independently across traders, so the panel reports marginals only.

Panel A: Hand-tuned model parameters

Symbol	Description	Value
N	number of traders	100,000
M	number of markets	50,000
σ_p	price-noise SD (log-odds)	0.5
α_W	mean log initial wealth	5.5
σ_W	log-wealth SD	1.5
α_τ	mean log signal precision	1.0
σ_τ	log-precision SD	1.5
α_μ	mean logit maker propensity	-2.0
σ_μ	logit-propensity SD	1.5
α_b	mean overconfidence wedge	2.0
σ_b	wedge SD	1.0
τ_0	prior precision	2.0
κ_0	trading aggressiveness	0.005
γ	precision exponent in κ_i	0.5
$\bar{\omega}$	max position fraction of wealth	0.05
α_A	participation intercept	-4.5
$\beta_{A,W}$	participation loading on log wealth	0.3
h_0	constant half-spread	0.005
$W_{\min}$	exit threshold	1.0

Panel B: Cross-sectional descriptive statistics by scenario

Scenario	$\log W$ mean	$\log W$ SD	$\log \tau$ mean	$\log \tau$ SD	logit μ mean	logit μ SD	b mean	b SD
Simulation Baseline	5.49	1.51	1.00	0.00	-2.00	0.00	0.00	0.00
+ Skilled Traders	5.49	1.51	1.00	1.50	-2.00	0.00	0.00	0.00
+ Maker Specialization	5.49	1.51	1.00	0.00	-2.00	1.49	0.00	0.00
+ Overconfidence	5.49	1.51	1.00	0.00	-2.00	0.00	2.00	1.00
+ All Three	5.49	1.51	1.00	1.50	-2.00	1.49	2.00	1.00

Figure I3: Simulated Price Calibration

This figure shows the simulated counterpart of Panel A of Figure 2 for each specification described in Section 7, anchored on the barebones *Simulation baseline* with each behavioral and informational channel turned on in isolation, then jointly. In each panel, markets are binned into ten equal-width price bins (0–10¢, 10–20¢, ..., 90¢–\$1). The x-axis is the bin midpoint price; the y-axis is the volume-weighted realized frequency at which the reference outcome resolves in-the-money. The dashed 45-degree line marks perfect calibration. The shaded band is a 95% bootstrap confidence interval (500 replications, resampled at the market level). The simulation uses $N = 100,000$ traders and $M = 50,000$ markets.




Figure I4: Simulated Volume-Weighted Excess Hit Rate by PnL Percentile

This figure shows the simulated counterpart of Figure 4 for each specification described in Section 7, anchored on the barebones *Simulation baseline* with each behavioral and informational channel turned on in isolation, then jointly. Each panel plots the mean (solid) and median (dashed) volume-weighted per-trade EHR by simulated PnL percentile; PnL percentile 0 corresponds to the least profitable users and 100 to the most profitable. The simulation uses $N = 100,000$ traders and $M = 50,000$ markets.





Figure I5: Simulated Maker Share by PnL Percentile

This figure plots the mean (solid) and median (dashed) maker share of volume by simulated PnL percentile, for each specification described in Section 7, anchored on the barebones *Simulation baseline* with each behavioral and informational channel turned on in isolation, then jointly. PnL percentile 0 corresponds to the least profitable users and 100 to the most profitable. Horizontal dotted lines in each panel mark the empirical maker-share targets reported in Table 5 (top 0.1% earners and bottom 95%). The simulation uses $N = 100,000$ traders and $M = 50,000$ markets.




J. Heterogeneity of Winners and Losers

This appendix supplements the main-text discussion of which traders win and lose by plotting the tail distributions of activity, liquidity provision, and price-region exposure for the top and bottom 100 users by terminal PnL. Section 5 compares means across PnL groups; the figures here show the heterogeneity behind those means at the tails.

Figure J6 plots a histogram for each of seven user-level features, with one series per group: the 100 users with the highest terminal PnL (green) and the 100 users with the lowest terminal PnL (red). Each series is a density histogram with 10 bins. For variables that span several orders of magnitude (number of trades, notional volume, number of markets), bins are log-spaced and the x -axis is on a log scale.

Figure J6: Distribution of User Features for Top 100 Winners and Losers

Each panel plots a density histogram (10 bins) of one user-level feature, with one series for the top 100 users by terminal mark-to-market PnL (green) and one for the bottom 100 (red). “Largest Market Share of Total Volume” is the share of a user’s total trade volume concentrated in their single most-traded market. For features that span several orders of magnitude, bins are log-spaced and the x -axis is on a log scale. The sample period is from November 11, 2022 to March 29, 2026.



Tables J22 and J23 list, respectively, the 100 users with the highest terminal mark-to-market PnL (the winners) and the 100 users with the lowest (the losers). Within each table, rows are ordered by absolute PnL in descending order. For each user we report the same seven characteristics shown in Figure J6: number of trades, notional volume in USDC, fraction of maker volume, fraction of trades at extreme prices ($p < 0.10$ or $p > 0.90$), number of markets traded, category HHI, and largest market share of total volume. We additionally report the user’s top category, defined as the market category whose summed per-market PnL contribution has the largest absolute value for that user; per-market PnL contributions are reconstructed from clean trades and the terminal mark-to-market on residual token positions, and markets without a category label are bucketed under *Untagged*. To preserve some privacy we identify each user only by the last six hex characters of the wallet address.

Table J22: Top 100 Users by Terminal PnL (Winners)

Note: The 100 users with the highest terminal mark-to-market PnL, ordered by absolute PnL descending. Wallets are shown by the last six hex characters of the address. The Top Category column is the market category whose summed per-market PnL contribution has the largest absolute value for that user; markets without a category in clean_markets are bucketed under Untagged.

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
1	f55839	22,029,972.64	19,695	43,013,259	0.4545	0.0004	14	1.0000	0.3516	Politics
2	6ad0cf	16,578,002.75	30,201	76,611,317	0.3883	0.1522	45	0.9980	0.3833	Politics
3	3033ee	10,409,176.82	162,330	265,514,610	0.7260	0.0410	2,303	0.9935	0.0075	Sports
4	3c6b76	8,705,078.46	6,649	16,402,745	0.3952	0.0000	7	1.0000	0.4415	Politics
5	0b0f29	7,698,903.29	18,074	40,551,791	0.2926	0.0005	8	0.6029	0.7268	Politics
6	efaa53	7,527,260.04	7,646	13,983,231	0.4809	0.0000	8	1.0000	0.4659	Politics
7	7875ea	6,470,439.58	1,749,437	332,437,290	0.8212	0.1329	44,810	1.0000	0.0024	Sports
8	d9f887	6,080,132.35	8,592	23,520,810	0.4581	0.0000	14	1.0000	0.5183	Politics
9	ee7091	5,936,332.02	5,793	32,196,692	0.7493	0.0585	1	1.0000	1.0000	Politics
10	0f4c7a	5,709,662.64	4,614	20,086,062	0.2897	0.0000	14	1.0000	0.3116	Sports
11	1ac270	5,634,963.92	7,563	11,212,840	0.5098	0.0000	7	1.0000	0.4724	Politics
12	a95e14	5,607,420.05	3,263,133	604,671,119	0.1826	0.0919	67,116	0.9977	0.0022	Sports
13	a35fcb	5,134,848.27	12,771	10,884,765	0.5863	0.0000	7	1.0000	0.4900	Politics
14	7d5da2	4,958,334.99	163,746	529,047,709	0.6375	0.2512	4,907	0.9681	0.0128	Sports

Continued on next page

Table J22 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
15	345565	4,800,670.78	3,395	9,949,069	0.4047	0.0000	6	1.0000	0.5744	Politics
16	88beae	4,723,848.38	1,222,657	295,611,242	0.7085	0.0480	46,877	1.0000	0.0017	Sports
17	5c16ea	4,552,611.03	5,788	60,070,363	0.8962	0.0271	428	0.9866	0.0204	Sports
18	6d7ab6	4,551,942.18	143,873	185,491,942	0.6507	0.0061	4,549	1.0000	0.0135	Sports
19	63be51	4,139,589.96	18,102	199,137,317	0.0516	0.0128	149	0.9975	0.0391	Sports
20	2299e3	4,042,384.90	7,508	9,457,725	0.4991	0.0000	11	1.0000	0.6015	Politics
21	a6c56e	3,807,135.56	267,852	225,617,089	0.8267	0.4662	1,458	1.0000	0.0239	Sports
22	08f9a2	3,741,724.97	2,768	13,750,267	0.3171	0.0000	8	1.0000	0.2759	Sports
23	579f3c	3,667,734.33	1,301	9,413,715	0.3112	0.0000	3	1.0000	0.7382	Sports
24	a8a881	3,644,213.61	133,032	192,646,025	0.5848	0.0018	6,336	0.9998	0.0116	Sports
25	b36356	3,430,433.35	13,337	15,831,477	0.6442	0.0063	24	0.9886	0.8390	Politics
26	5aa429	3,402,488.73	50,404	82,906,349	0.5848	0.0059	802	1.0000	0.0580	Sports
27	5c5f6d	3,132,785.80	6,955	17,303,666	0.2593	0.1145	121	0.8322	0.1388	Sports
28	a30d50	3,114,400.67	3,999	8,794,465	0.1871	0.0010	3	1.0000	0.8426	Politics
29	e3debf	3,094,564.50	1,149,638	347,721,723	0.7453	0.0262	36,648	1.0000	0.0056	Sports
30	b03dd0	3,092,634.86	6,032	8,450,963	0.4408	0.0008	13	1.0000	0.6961	Politics
31	b35029	2,625,243.96	9,089	32,510,191	0.0420	0.0058	561	1.0000	0.0185	Sports
32	91456e	2,604,488.90	1,291	8,132,716	0.1578	0.0000	33	1.0000	0.1109	Sports

Continued on next page

Table J22 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
33	af1697	2,532,416.28	81,301	81,982,674	0.6110	0.0396	571	0.9969	0.0114	Sports
34	dc1344	2,515,938.77	625,540	470,350,236	0.7357	0.4715	9,550	0.6870	0.0425	Politics
35	7a9723	2,411,291.74	30,418	33,020,925	0.5096	0.1080	660	0.9569	0.0115	Sports
36	f9ba9a	2,376,631.38	6,841,716	199,693,236	0.6125	0.0347	33,951	1.0000	0.0010	Crypto
37	10c684	2,358,105.09	74,209	81,353,082	0.8104	0.1007	601	1.0000	0.0175	Sports
38	9cb18d	2,322,418.55	150,107	98,142,942	0.8141	0.0490	9,441	0.9999	0.0039	Sports
39	9afbd7	2,309,156.74	3,028	14,869,434	0.3532	0.0000	8	1.0000	0.5749	Sports
40	d6ffbe	2,262,965.14	197,996	336,460,069	0.6127	0.1686	659	0.7631	0.3724	Politics
41	f87e44	2,243,149.16	180,872	97,564,244	0.4845	0.2025	990	0.9185	0.0461	Sports
42	c5cea1	2,212,952.49	103,145	164,414,911	0.6522	0.5733	7,064	0.9562	0.0067	Sports
43	14e33d	2,210,647.43	40,698	120,653,616	0.7529	0.0003	2,574	0.9918	0.0099	Sports
44	51141d	2,203,505.00	26,071	53,213,991	0.7031	0.4592	1,341	0.6705	0.0280	Sports
45	fb9bc1	2,147,901.76	104,588	246,428,072	0.6253	0.1986	1,249	0.9963	0.0239	Sports
46	334d20	2,097,290.34	1,501	5,237,151	0.1327	0.0000	9	1.0000	0.4446	Sports
47	467cd0	2,092,886.41	1,393	3,655,047	0.2768	0.0000	6	1.0000	0.5878	Politics
48	13ebc1	1,995,476.57	86,346	200,944,880	0.7199	0.5729	2,436	0.5114	0.0299	Politics
49	598ff7	1,990,361.97	5,664	22,184,124	0.2808	0.0000	20	1.0000	0.1808	Sports
50	f5594f	1,896,440.47	11,148	5,182,301	0.6277	0.1023	19	0.9977	0.5791	Politics

Continued on next page

Table J22 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
51	453ca7	1,878,805.19	83,942	61,841,052	0.0463	0.5937	1,326	1.0000	0.0150	Crypto
52	a53505	1,849,163.67	2,623	7,016,597	0.6476	0.0038	9	1.0000	0.7194	Politics
53	1289e4	1,771,355.29	9,475	14,530,562	0.7167	0.5187	66	0.8452	0.4900	Politics
54	d5386e	1,766,564.93	920	4,549,553	0.0013	0.0022	8	1.0000	0.6401	Politics
55	b504d0	1,738,460.04	1,389	6,769,192	0.1608	0.3708	20	0.6903	0.3549	Politics
56	0993aa	1,714,234.58	10,009,113	208,947,314	0.5852	0.0615	54,356	1.0000	0.0008	Crypto
57	974222	1,687,089.53	192,002	75,695,854	0.7600	0.0637	11,041	0.9962	0.0067	Sports
58	3102c3	1,647,073.58	12,232	11,839,832	0.5300	0.4362	56	0.5981	0.2517	Sports
59	755009	1,579,123.61	205,909	90,777,952	0.3084	0.0810	11,435	0.9634	0.0121	Sports
60	7dd2e5	1,578,604.37	61,235	170,863,143	0.7001	0.9173	1,281	0.4493	0.0472	Politics
61	a2604a	1,554,147.98	10,113	12,196,466	0.5838	0.0000	76	1.0000	0.0688	Sports
62	644c82	1,502,176.65	377,634	97,783,815	0.5742	0.1371	14,879	0.9031	0.0121	Crypto
63	0a758e	1,497,726.49	47,318	63,507,521	0.3187	0.6407	1,088	0.6594	0.0288	Politics
64	b47a07	1,496,247.79	6,786	48,318,877	0.0840	0.0024	120	1.0000	0.0281	Sports
65	aba3bd	1,475,486.22	5,746	28,493,252	0.2335	0.0010	109	1.0000	0.0348	Sports
66	8ffed6	1,473,195.49	63,053	73,147,942	0.3733	0.0083	845	1.0000	0.0153	Sports
67	24f59d	1,456,737.32	2,151	8,858,267	0.2412	0.3854	30	0.4159	0.3911	Politics
68	cd5817	1,443,591.16	85,150	60,385,427	0.5452	0.2943	1,906	0.3613	0.0191	Politics

Continued on next page

Table J22 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
69	9f0977	1,435,937.42	71,416	158,508,752	0.4838	0.3438	1,944	0.9873	0.0203	Sports
70	52aeeb	1,419,783.36	62,008	58,307,104	0.8366	0.0103	1,578	0.9998	0.0099	Sports
71	96d3e7	1,405,950.09	8,190	14,655,550	0.6728	0.3370	125	0.4933	0.1143	Politics
72	9d08d5	1,345,991.49	41,772	64,368,306	0.5496	0.3152	1,281	0.4884	0.0668	Politics
73	c1a83c	1,308,310.55	3,905	7,213,168	0.7811	0.2927	150	0.6694	0.4507	Politics
74	65e113	1,291,697.44	1,804	10,820,687	0.5633	0.0316	29	1.0000	0.1164	Sports
75	69a8db	1,264,492.87	38,467	46,644,930	0.0393	0.0173	3,843	1.0000	0.0078	Sports
76	44c255	1,252,107.95	5,423	35,852,301	0.8930	0.0380	417	1.0000	0.0370	Sports
77	e823e1	1,223,625.06	877,318	407,033,887	0.8583	0.8261	3,977	0.3935	0.0259	Politics
78	239a09	1,202,903.12	1,687	6,488,687	0.4870	0.0000	41	1.0000	0.0925	Sports
79	18dae8	1,189,727.32	2,950	23,095,757	0.1697	0.0000	65	1.0000	0.0509	Sports
80	b0ab35	1,178,169.97	269,441	102,391,326	0.5189	0.4790	2,492	0.6913	0.0287	Politics
81	d2dbbd	1,170,897.20	7,303	30,428,338	0.3934	0.0800	111	1.0000	0.0367	Sports
82	1ea5e5	1,170,581.62	522	3,170,054	0.0000	0.0000	1	1.0000	1.0000	Politics
83	308057	1,140,869.31	7,505	12,663,462	0.2160	0.0452	59	1.0000	0.0462	Sports
84	aae842	1,140,704.53	2,709	6,638,979	0.0143	0.0078	95	1.0000	0.5139	Sports
85	637c9d	1,135,012.89	3,892,050	122,060,203	0.4478	0.1423	25,123	1.0000	0.0010	Crypto
86	1e5c6b	1,095,017.51	282,522	144,055,460	0.6605	0.3883	6,446	0.5866	0.0230	Politics

Continued on next page

Table J22 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
87	5c32f2	1,092,444.13	67,660	47,404,791	0.6300	0.0289	1,473	1.0000	0.0197	Sports
88	cf87b3	1,082,825.09	278,542	172,891,375	0.6335	0.4895	3,600	0.8108	0.0892	Politics
89	0079a0	1,081,855.33	1,159	6,660,654	0.5405	0.0026	13	1.0000	0.1847	Sports
90	d200f2	1,081,137.58	15,328	20,841,865	0.5876	0.2328	286	0.4267	0.0766	Politics
91	4a752c	1,080,542.78	41,235	33,852,796	0.0881	0.0224	1,439	0.9357	0.0085	Sports
92	8d7422	1,069,031.05	5,257	14,225,930	0.2219	0.0010	178	1.0000	0.0196	Sports
93	ba9c95	1,061,379.57	28,716	78,565,174	0.7939	0.1579	922	0.5444	0.0239	Sports
94	1f43bc	1,029,448.97	1,689	6,498,120	0.0294	0.0018	73	0.9521	0.0886	Sports
95	e70853	1,025,563.04	61,316	33,428,471	0.5677	0.1401	2,895	0.9462	0.0105	Sports
96	df4429	1,023,723.24	43,311	44,102,106	0.4557	0.0994	484	0.9717	0.0123	Sports
97	a253ba	1,020,255.47	10,734	7,280,730	0.6735	0.0000	68	1.0000	0.1307	Sports
98	fd1292	1,019,538.20	29,723	39,993,246	0.5176	0.3907	1,555	0.3994	0.0380	Politics
99	5beed6	1,015,348.71	11,025	8,226,581	0.4698	0.5547	48	0.9994	0.2650	Politics
100	74f05d	1,005,286.13	44,781	45,729,882	0.6136	0.0246	3,301	0.9906	0.0046	Sports

Table J23: Top 100 Users by Terminal PnL (Losers)

Note: The 100 users with the lowest terminal mark-to-market PnL, ordered by absolute PnL descending. Wallets are shown by the last six hex characters of the address. The Top Category column is the market category whose summed per-market PnL contribution has the largest absolute value for that user; markets without a category in clean_markets are bucketed under Untagged.

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
1	38dd4c	-10,022,403.27	19,531	150,299,249	0.3114	0.0000	381	1.0000	0.0117	Sports
2	8a3782	-7,172,406.18	51,466	434,763,592	0.0140	0.0024	1,399	1.0000	0.0028	Sports
3	f5cc34	-6,799,218.37	1,254	19,066,349	0.6985	0.6021	5	1.0000	0.6931	Politics
4	6bd96a	-5,521,693.91	62,537	122,896,778	0.6347	0.0508	835	1.0000	0.0245	Sports
5	3bc735	-5,031,568.42	2,708	13,376,552	0.3649	0.0000	11	1.0000	0.3852	Sports
6	046883	-5,019,314.14	1,493	13,499,020	0.7975	0.0884	1	1.0000	1.0000	Politics
7	a86825	-4,814,564.34	120,669	62,093,782	0.9216	0.3506	382	0.8832	0.0337	Politics
8	53dd8d	-4,668,525.41	6,503	20,065,920	0.5681	0.2194	39	0.9952	0.4429	Politics
9	8b259e	-4,535,121.70	10,987	37,206,907	0.2941	0.0234	85	1.0000	0.0331	Sports
10	6d5408	-3,820,163.48	10,093	49,919,524	0.3392	0.0229	191	1.0000	0.0275	Sports
11	ee32e8	-3,738,937.27	7,815	10,211,476	0.8750	0.3152	20	1.0000	0.7227	Politics
12	1df3c4	-3,677,853.43	29,230	144,256,999	0.3193	0.0194	334	1.0000	0.0173	Sports
13	c2aeaa	-3,332,045.94	347,922	130,405,190	0.0171	0.1434	4,574	0.9674	0.0313	Sports
14	accc57	-3,238,246.38	3,294	10,991,682	0.2727	0.0000	10	1.0000	0.3435	Sports

Continued on next page

Table J23 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
15	fb33fb	-2,979,234.47	1,652	5,429,822	0.5279	0.0000	2	1.0000	0.9278	Sports
16	b4a3cd	-2,942,208.83	44,658	55,838,279	0.5531	0.0186	402	1.0000	0.0318	Sports
17	a47ebb	-2,934,355.50	1,739	8,941,714	0.0877	0.0811	22	0.7250	0.2432	Sports
18	d0eec8	-2,782,416.13	104,050	129,411,154	0.8341	0.1128	3,234	0.6895	0.0116	Sports
19	0ebd89	-2,567,727.50	1,262	6,563,349	0.0603	0.0151	10	1.0000	0.6864	Politics
20	d4340f	-2,528,996.19	147	3,696,321	0.0827	0.0000	1	1.0000	1.0000	Sports
21	f9f9be	-2,477,417.73	2,402	10,530,560	0.4854	0.0000	24	0.9999	0.1579	Sports
22	485210	-2,439,303.89	75,035	192,273,770	0.6324	0.0133	508	1.0000	0.0195	Sports
23	e976bf	-2,368,502.90	4,981	8,942,645	0.2461	0.4340	121	0.9445	0.3060	Politics
24	e6af98	-2,278,707.72	4,613	16,263,076	0.4880	0.0570	38	1.0000	0.0686	Sports
25	67c46b	-2,231,464.40	40,391	102,811,036	0.6607	0.0228	263	1.0000	0.0355	Sports
26	3e3f84	-2,224,397.87	67,400	79,345,211	0.5650	0.5711	1,128	0.5421	0.2358	Politics
27	a374ce	-2,118,209.94	2,503	11,540,727	0.5470	0.0935	48	0.7490	0.5431	Politics
28	d038f9	-2,094,891.17	1,401	5,263,889	0.4289	0.0000	3	1.0000	0.7470	Politics
29	bc3532	-2,044,654.24	53,916	57,984,013	0.5026	0.0056	2,275	0.9972	0.0060	Sports
30	854ebe	-2,041,618.74	6,590	34,099,634	0.8900	0.0938	231	0.8467	0.0642	Sports
31	f4d49e	-2,040,886.72	433	7,089,338	0.1614	0.0000	6	1.0000	0.2565	Sports
32	d8a7fd	-2,002,352.44	806	6,292,417	0.0234	0.2655	19	1.0000	0.6941	Politics

Continued on next page

Table J23 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
33	850782	-1,999,445.56	1,730	5,124,000	0.1535	0.0000	1	1.0000	1.0000	Politics
34	f58ac6	-1,999,122.03	454	5,405,172	0.0312	0.0000	1	1.0000	1.0000	Politics
35	628b48	-1,998,510.08	1,721	5,201,376	0.7775	0.0000	2	1.0000	0.5306	Politics
36	76f52c	-1,927,056.05	22,543	121,784,864	0.4777	0.3398	111	0.6073	0.1517	Politics
37	166b14	-1,892,087.97	13,813	52,374,659	0.3041	0.0164	146	1.0000	0.0402	Sports
38	a053be	-1,837,916.52	3,859	17,475,404	0.4992	0.0417	85	1.0000	0.0594	Sports
39	da2813	-1,785,377.94	2,180	15,037,808	0.1932	0.0009	30	1.0000	0.0832	Sports
40	c9bba3	-1,767,964.25	17,107	36,686,600	0.4141	0.1500	541	1.0000	0.0437	Sports
41	d18df0	-1,640,857.26	2,266	3,071,959	0.7164	0.0000	2	1.0000	0.5155	Politics
42	30e0de	-1,621,408.13	5,247	17,541,721	0.5531	0.0046	61	1.0000	0.0522	Sports
43	1004e8	-1,510,146.76	5,431	14,482,635	0.5591	0.0897	162	0.5294	0.1495	Politics
44	554513	-1,505,350.25	5,103	4,027,852	0.6298	0.0000	2	1.0000	0.5303	Politics
45	cd6ef7	-1,505,164.49	2,302	7,270,595	0.1479	0.1994	42	0.8703	0.3892	Politics
46	80f171	-1,449,619.73	1,484	3,773,102	0.3442	0.0000	9	1.0000	0.2944	Sports
47	ab39ae	-1,445,250.18	7,230	57,107,550	0.0852	0.0820	130	1.0000	0.0456	Sports
48	d4d1a2	-1,442,846.72	196,425	46,349,191	0.0507	0.2339	1,279	0.8602	0.0530	Sports
49	4974cc	-1,424,796.76	10,242	48,840,485	0.6131	0.0158	333	1.0000	0.0300	Sports
50	42e5b3	-1,397,704.50	1,023	11,897,153	0.4231	0.0127	6	0.9708	0.6415	Politics

Continued on next page

Table J23 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
51	f51fe1	-1,397,086.55	2,253	7,707,365	0.4048	0.0000	6	1.0000	0.3545	Sports
52	a7cd31	-1,307,379.40	5,373	9,613,384	0.6112	0.0350	37	1.0000	0.0552	Sports
53	871ef0	-1,286,789.57	15,261	10,327,228	0.5396	0.0000	90	1.0000	0.0541	Sports
54	8a1e4a	-1,279,230.55	3,333	20,479,244	0.2025	0.0177	94	1.0000	0.0526	Sports
55	a328ad	-1,273,709.66	23,263	32,337,796	0.0000	0.0484	643	0.9939	0.0105	Sports
56	1372d6	-1,258,742.23	587	2,604,580	0.6401	0.1022	6	1.0000	0.4188	Politics
57	ae6a47	-1,231,673.85	1,035,649	150,698,086	0.0782	0.0325	15,120	0.9973	0.0093	Sports
58	276f2c	-1,225,532.27	4,673	7,107,980	0.1254	0.0460	1	1.0000	1.0000	Politics
59	b30f7b	-1,214,111.41	4,207	36,717,035	0.0561	0.0863	44	1.0000	0.0791	Sports
60	2ae5cd	-1,200,376.73	75,964	68,355,519	0.0398	0.0018	7,126	1.0000	0.0048	Sports
61	21b0cf	-1,195,531.44	8,890	17,270,341	0.6400	0.4744	93	0.9687	0.2155	Politics
62	cd5133	-1,179,947.08	3,462	3,493,660	0.8734	0.2441	21	0.8887	0.2985	Politics
63	26d3a8	-1,178,829.94	2,299	4,695,519	0.7257	0.0687	22	0.6239	0.3922	Politics
64	30c76a	-1,134,335.03	858,804	159,886,969	0.0664	0.0261	11,153	0.9777	0.0059	Sports
65	82d698	-1,127,287.82	7,816	19,166,687	0.6634	0.0306	201	1.0000	0.0162	Sports
66	241d0a	-1,120,145.86	5,041	14,880,253	0.2241	0.1285	59	0.7509	0.1952	Politics
67	fad24d	-1,110,281.28	51,948	72,442,900	0.6240	0.5506	1,759	0.7731	0.0240	Sports
68	19fa74	-1,089,985.00	2,382	6,725,571	0.3409	0.0000	9	1.0000	0.2178	Sports

Continued on next page

Table J23 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
69	60a818	-1,087,249.06	2,463	5,661,908	0.7321	0.3979	9	1.0000	0.4500	Politics
70	e1fd3d	-1,064,490.01	1,887	2,811,819	0.5167	0.0000	7	1.0000	0.4593	Sports
71	6a09cb	-1,058,855.84	20,062	61,680,625	0.8104	0.0047	719	0.9696	0.0138	Sports
72	84a7a1	-1,049,152.40	927	7,832,698	0.0273	0.0000	23	1.0000	0.2545	Sports
73	770af8	-1,048,655.09	583	5,059,266	0.7261	0.0000	13	1.0000	0.2950	Sports
74	006f95	-1,039,572.88	20,969	13,466,206	0.5470	0.0008	148	1.0000	0.0439	Sports
75	43b65f	-1,037,202.58	378	1,825,162	0.7815	0.0582	7	1.0000	0.8126	Politics
76	b2bff2	-1,035,307.46	1,614	3,108,111	0.2459	0.0000	5	1.0000	0.6079	Sports
77	cc6925	-1,026,369.58	3,186	3,876,930	0.9141	0.0512	6	1.0000	0.3983	Politics
78	8a6beb	-1,014,719.87	2,259	13,370,312	0.2295	0.0000	9	1.0000	0.4003	Sports
79	a15a96	-1,010,198.84	682	1,957,609	0.8527	0.0000	1	1.0000	1.0000	Politics
80	e561e1	-999,752.37	1,211	2,126,667	0.4314	0.0000	1	1.0000	1.0000	Politics
81	acd77b	-997,412.97	180	2,206,356	0.0000	0.0000	1	1.0000	1.0000	Politics
82	df9c69	-994,711.15	699	2,652,020	0.5885	0.0000	2	1.0000	0.9993	Politics
83	fc15bd	-986,935.21	11,349	12,423,262	0.6881	0.0859	161	0.9591	0.0344	Sports
84	2150db	-983,182.89	1,027	2,804,862	0.0744	0.0000	1	1.0000	1.0000	Politics
85	1f69d9	-982,282.07	676	6,674,617	0.3616	0.0695	4	1.0000	0.8449	Politics
86	696db6	-974,933.37	2,357	2,336,328	0.4103	0.0467	33	0.9874	0.6134	Politics

Continued on next page

Table J23 – continued

Rank	Wallet (last 6)	PnL (USDC)	Trades	Volume (USDC)	Frac Maker Vol.	Frac Extreme Price	Markets	Cat. HHI	Max Mkt Share	Top Category
87	30a59b	-974,826.29	26,777	29,574,295	0.5752	0.0247	2,697	1.0000	0.0250	Sports
88	67817e	-972,901.16	5,358	23,674,706	0.2631	0.3044	160	0.6360	0.3401	Politics
89	b7b731	-941,684.93	1,218	4,656,157	0.7345	0.0435	6	1.0000	0.6158	Politics
90	d444a5	-931,048.13	9,125	24,550,708	0.1593	0.2613	178	0.6134	0.1639	Politics
91	22aaaae	-920,686.73	303	2,034,597	0.1638	0.0000	3	1.0000	0.5164	Sports
92	53b215	-919,315.12	6,252	8,524,046	0.3618	0.0198	46	1.0000	0.0817	Sports
93	765055	-908,277.42	20,859	18,036,495	0.4517	0.1193	464	0.2439	0.0564	Politics
94	f7d196	-903,643.87	35,443	21,318,813	0.3147	0.0049	404	1.0000	0.0304	Sports
95	30364a	-899,971.65	236	2,444,892	1.0000	0.0000	1	1.0000	1.0000	Politics
96	bdfa77	-897,726.05	3,775	9,768,504	0.2701	0.4053	111	0.8279	0.0972	Sports
97	8f6d03	-896,968.62	467	1,804,387	0.8992	0.0343	4	1.0000	0.5177	Politics
98	b78240	-894,331.39	1,227	3,369,052	0.0305	0.0000	5	0.9813	0.9302	Politics
99	0574fe	-893,852.67	3,968	13,712,670	0.1861	0.0028	171	1.0000	0.0516	Sports
100	b44908	-892,543.69	990	9,061,856	0.3977	0.0354	45	1.0000	0.1070	Sports