

Manufacturing Speculation: Contract Design and Trader Performance in Prediction Markets

Yinru Lei *
Hunan University

Alberto G. Rossi †
Georgetown University

September 14, 2026

Abstract

We study how the contracts a prediction market chooses to offer shape trading activity and trader performance. Using the complete trading record of Polymarket’s financial contracts from February 2024 to March 2026, we document a shift from one-off and monthly contracts to five-minute bets on whether asset prices would rise or fall. Trading rapidly followed each new design, increasing trading intensity as trader performance deteriorated. The share of profitable wallets fell from 60% in monthly contracts to 25% in five-minute contracts. At observed loss rates and trading intensity, \$100 of capital would be depleted in 404 days in monthly contracts, compared with just 1.6 days in five-minute contracts. New accounts sustained volume as existing traders left. We develop a model that explains this pattern: a platform seeking to maximize trading volume offers shorter contracts because each new contract renews opportunities to bet on disagreement. Repeated betting then erodes the wealth of most traders, even without fees. Our findings reveal a tension between platform growth and trader performance: absent regulatory interventions, contract designs that encourage more trading can also accelerate trader losses.

Keywords: Security Design, Financial Innovation, Prediction Markets, Speculation, Retail Investors, Polymarket

JEL Classification: D47, G13, G14, G18, G41

*College of Finance and Statistics, Hunan University. Email: yinrulei@hnu.edu.cn.

†McDonough School of Business, Georgetown University. Email: agr60@georgetown.edu.

1 Introduction

A large literature studies prediction markets as devices for aggregating dispersed information (Wolfers and Zitzewitz, 2004; Manski, 2006; Snowberg, Wolfers, and Zitzewitz, 2007), and their prices are increasingly used as measures of macroeconomic expectations (Diercks, Katz, and Wright, 2026; Angelini, 2026), of corporate earnings expectations (Gomez Cram et al., 2025; Rabetti, Shao, and Zhang, 2026; Glaze, López-Vilaró, and Stephan, 2026), of central-bank credibility (Eichengreen et al., 2025), and of price discovery across venues (Ng et al., 2026; Fabi et al., 2026). This literature takes as given the set of contracts from which the prices arise. Which events become tradable, in what form, and at what maturity are, however, choices of the platform that lists the contracts, and the objective those choices serve need not coincide with the informational role that makes prediction markets valuable to researchers. We study these listing choices and their consequences for the traders who adopt the resulting contracts.

The theory of security design offers a benchmark for how a self-interested venue makes these choices. Duffie and Jackson (1989) show that an exchange that maximizes trading volume lists the contract that generates the largest positions, and that this contract need not be the one that raises traders’ welfare. Simsek (2013b) shows that when investors disagree, the assets a profit-seeking intermediary introduces are those that facilitate betting on the disagreement, raising speculative risk rather than improving risk sharing; Gennaioli, Shleifer, and Vishny (2012) obtain a related distortion when investors neglect tail risks. Empirically, however, the design problem is difficult to observe. Studies of retail structured products and specialized funds find that issuers cater to investor demand at investors’ expense (Célérier and Vallée, 2017; Vokata, 2021; Ben-David et al., 2023), but they observe heterogeneous products issued by many intermediaries rather than the sequence of design choices of a single venue, and they observe flows at the product level rather than the investors who adopt each design and their subsequent performance.

Our setting is Polymarket, the largest blockchain-based prediction market. We assemble the complete record of the platform’s trading activity from its public data interfaces—a ledger of roughly one billion transaction records—and study the 221,000 resolved markets written on financial underlying assets and the 1.25 million wallets that trade them from February 2024, when the platform’s financial

series become active, through March 2026. The setting offers two advantages for studying security design. The first is that the platform’s listing decisions were not reviewed by any regulator. On regulated U.S. derivatives exchanges, a new contract is listed upon the exchange’s self-certification to the CFTC and remains subject to the Commission’s review ([Lockhart, 2025](#); [Bartlett and O’Hara, 2026](#)). Polymarket operated outside this apparatus for our entire sample period: barred from serving U.S. customers by a 2022 CFTC settlement, the global exchange ran offshore, without registration, without ex ante review of its listings, and—for most of the period—without trading fees. The regulated U.S. venue that the company launched near the end of our sample is a legally separate exchange and does not enter our data; every listing decision we study, including the introduction of five-minute contracts in February 2026, was the platform’s own.¹ The menu we study is therefore a record of what a trading platform chooses to create when no regulator reviews its listings. The second advantage is that both sides of the design problem are observable: the platform repeatedly lists contracts on the same underlying assets while varying their payoff structure and issuance frequency, and wallet-level histories reveal which accounts adopt each design, where they were previously active, and how they subsequently perform.

To interpret the platform’s choices, we develop a model in which a volume-maximizing platform designs the contracts through which traders with heterogeneous beliefs speculate. The trader side follows [Simsek \(2013b\)](#): positions decompose into a hedging component and a speculative component that grows with disagreement. The platform side follows [Duffie and Jackson \(1989\)](#): the platform selects the contract design that maximizes trading volume. The model yields four results. First, the volume-maximizing contract concentrates its payoff on the outcome about which traders disagree. Second, the same objective compresses the contract’s horizon, because shorter contracts renew disagreement—and the trading it generates—more frequently. Third, repetition is costly for traders: because wealth compounds multiplicatively, repeated fair bets carry a strictly negative geometric drift, so the majority of traders lose even in the absence of fees, and losses accumulate faster at shorter horizons. Fourth, as

¹In January 2022, the CFTC ordered Polymarket to cease offering unregistered binary options to U.S. customers, and the global exchange thereafter operated offshore. In November 2025, the CFTC amended the order of designation of QCX LLC, the licensed exchange the company had acquired in July 2025, to permit intermediated trading, and the company launched a separate U.S. platform under that designation in December 2025; the global exchange sought approval to admit U.S. traders only in April 2026, after our sample ends. In June 2026, the CFTC proposed a new regulatory framework for event contracts.

a matter of accounting, aggregate volume can be sustained only if new capital replaces what winning traders withdraw and losing traders deplete, so the platform’s growth depends on the recruitment of new accounts.

Guided by the model, we begin by documenting the evolution of the platform’s menu using 221,000 resolved markets written on financial underlyings. Between February 2024 and March 2026, Polymarket introduced five contract structures and eight issuance frequencies in strict sequence, and the sequence converged on the model’s volume-maximizing design. Payoff structures moved from individually specified thematic questions, through threshold and multi-threshold designs graded in the size of the price move, to the Up/Down contract—a single binary bet on the direction of the underlying price; issuance frequencies shortened step by step from one-off listings to monthly, weekly, daily, intraday, and ultimately five-minute contracts. New designs added to the menu rather than replacing it, but new listings concentrated in the fastest designs: the fifteen-minute and five-minute contracts alone account for roughly two thirds of all cryptocurrency markets the platform has ever listed, and the platform writes nearly 8,000 markets on the average crypto underlying. The two margins interact in the manner the model predicts: the richer structures never reach the fastest frequencies, and at horizons of fifteen minutes and shorter the binary Up/Down contract is the only design listed. Cryptocurrency leads the sequence, and equities and commodities follow it with a lag of approximately one year.

We then examine how trading activity responded to these introductions. Each newly introduced design absorbed the majority of trading volume within weeks of its launch: hit-price contracts captured 90% of crypto volume within a month of their introduction, Up/Down contracts approached 90% by the end of the sample, and five-minute contracts drew nearly two thirds of weekly volume within seven weeks of becoming active. Under the completed menu, Up/Down contracts account for 91% of crypto trading volume and the two fastest frequencies for 80%. Adoption operates through trading intensity rather than through the size of stakes: as contract duration falls by a factor of roughly 8,600 from monthly to five-minute contracts, the average stake falls only from \$99 to \$11, while the average number of trades per user per month rises from 5 to 343 and monthly volume per user rises from roughly \$500 to \$3,800. In the aggregate, volume rose by a factor of roughly 16,000 over the sample while the number of active traders rose by a factor of roughly 3,000, so the platform’s growth was

carried disproportionately by rising volume per trader.

Finally, we measure the consequences for the traders. Realized performance deteriorates along both design margins. The share of wallets ending in profit falls at every increase in issuance frequency—69% in one-off contracts, 60% in monthly, 50% in daily, 35% in fifteen-minute, and 25% in five-minute contracts—and mean returns fall from -1.9% to -22.5% along the same path. Estimates with wallet fixed effects indicate that the deterioration is not an artifact of which traders select into each design: within the same wallet, returns decline from approximately zero in monthly contracts to -16% at sub-hourly frequencies. Losses per position and trading intensity compound: from one-off to five-minute contracts the mean loss per position roughly quadruples while the number of positions per active day rises sixteen-fold, so \$100 of trading capital that is expected to last 404 days in monthly contracts is expected to last 1.6 days in five-minute contracts.

Retention follows the same pattern: 79% of wallets entering the platform through one-off contracts remain active after thirty days, against 12% of those entering through five-minute contracts. The platform’s volume is nonetheless sustained, because departing traders are replaced. By the end of the sample, roughly half of weekly volume is traded by accounts less than one month old—the continuing inflow that the model’s capital-flow condition requires. New accounts enter overwhelmingly through intermediate designs—weekly and hourly contracts receive three quarters of first-trade volume, and the two fastest designs only 15%—and subsequent positions migrate toward faster contracts almost without reversal, so the fastest designs are populated by internal migration rather than by direct entry.

Two features of the setting sharpen the interpretation of these results. Because Polymarket charged no explicit fees for most of the sample period, the deterioration of trader performance cannot be attributed to fees. Traders still cross a bid–ask spread, and part of what they lose accrues to the accounts that provide liquidity or enter earlier ([Akey et al., 2026](#); [Della Vedova, 2026b](#); [Gomez Cram et al., 2026](#)). The model’s volatility-drag mechanism shows why the losses require no transfer at all: repeated fair bets erode the wealth of most traders even when aggregate trading profits are zero. The absence of fees does not remove the platform’s stake in volume: volume is the headline measure by which prediction-market platforms are compared, and it became the base for explicit taker fees once the platform introduced them on its fastest cryptocurrency designs in the last months of the sample.

The second feature is the sample’s restriction to contracts written on financial underlying assets. For event markets—elections, sports outcomes, or geopolitical developments—the arrival of new trading opportunities is inherently tied to the arrival of new events in the world. Changes in the contract menu therefore combine variation in platform design with variation in the underlying events available to trade. For financial underlyings, the assets themselves trade continuously and are always available; what varies is the platform’s choice of how to package them. The menu variation we exploit is therefore attributable to the platform’s own listing decisions rather than to the arrival of outside events. Our claims accordingly concern the platform’s financial segment rather than the platform in its entirety.

The pattern we document parallels the concentration of retail speculation in the shortest-dated instruments on offer, most visible in equity options, where half of retail trades are in contracts with less than a week to expiration ([Bryzgalova, Pavlova, and Sikorskaya, 2023](#)) and same-day contracts have become a large share of index-option volume ([Beckmeyer, Branger, and Gayda, 2023](#)). Two findings distinguish our setting. First, we observe the supply side directly: the shortening of maturities on Polymarket is a sequence of listing decisions by a single platform, not the aggregation of many issuers’ responses to demand. Second, the population of traders adjusts through replacement rather than learning: the fastest designs are populated by a continually renewed inflow of new accounts rather than by incumbents who improve.

The results identify a margin for investor protection that operates before any trade occurs. Regulatory scrutiny of prediction markets has concentrated on trading conduct within listed contracts—insider trading ([Mitts and Ofir, 2026](#); [Liu, 2026](#)) and settlement manipulation ([Dai, Jia, and Yu, 2026](#))—and on which categories of event may be listed at all; the review frameworks under discussion evaluate contracts one at a time ([Lockhart, 2025](#)). The deterioration in trader outcomes documented here arises instead from the composition and speed of the menu as a whole, on a platform where each individual contract is small, transparent, and, for most of the sample, traded without fees. More broadly, the setting shows financial innovation proceeding under market forces alone, and the outcome suggests that market discipline by itself need not align innovation with investor welfare. Restricting the scale of retail speculation has improved retail outcomes elsewhere: a cap on retail forex leverage cut trading volume and the losses of the most leveraged traders without impairing liquidity ([Heimer](#)

and Simsek, 2019). The contract menu is the analogous margin in our setting: it determines how often traders can stake their capital and, as we show, how fast that capital is depleted.

The remainder of the paper proceeds as follows. Section 2 relates our findings to the literature. Section 3 describes the setting, the data, and the contract taxonomy. Section 4 develops the model, Section 5 presents the empirical analysis, and Section 6 concludes.

2 Related Literature

First, we contribute to the literature on security design and financial innovation. Duffie and Jackson (1989) characterize the futures contract that a volume-maximizing exchange lists, and show that the exchange can prefer a contract that raises turnover even when it lowers traders' expected utility. A related literature examines the innovation of individual futures contracts and the determinants of their success (Silber, 1981; Johnston and McConnell, 1989; Cuny, 1993; Tashjian and Weissman, 1995). In modern equity markets, Budish, Lee, and Shim (2024) show that competition among exchanges need not undo trader-harming designs, because exchange rents derive from the frictions those designs create. Simsek (2013b) shows that when investors disagree, new assets raise speculative risk faster than they improve risk sharing, and that a profit-seeking intermediary introduces the assets on which disagreement is easiest to trade; Simsek (2013a) shows that the same logic applies when trading itself becomes cheaper and faster: lowering the frictions that suppress bets raises speculative risk taking rather than risk sharing. Gennaioli, Shleifer, and Vishny (2012) provide a complementary mechanism: profit-seeking intermediaries cater to investor demand by engineering new securities, and when investors neglect downside risks, such innovation can benefit intermediaries while leaving investors exposed to risks they fail to anticipate. In our model, a platform with the volume objective of Duffie and Jackson (1989) designs contracts for the disagreeing traders of Simsek (2013b), choosing both what the contracts pay off on and how often they are reissued. Our paper therefore contributes to this literature by making the maturity a design variable: shorter maturities multiply the rounds of trading, and the repetition erodes the wealth of the traders who adopt the contracts.

On the empirical side, a security-design literature documents that issuers design retail products for the demand they can stimulate rather than the risks they can share. C  lerier and Vall  e (2017)

show that banks pair salient headline rates with complexity, and that the most complex structured products carry the highest markups and the worst ex post performance; [Vokata \(2021\)](#) finds that many yield-enhancement products are dominated state by state by portfolios of listed options; and [Ben-David et al. \(2023\)](#) show that specialized ETFs launched on attention-grabbing themes attract flows and lose roughly 30% risk-adjusted over their first five years. We extend this evidence on both sides of the market. On the supply side, these studies observe products that differ across issuers, underlying assets, and issuance terms; in our setting a single platform repeatedly lists contracts on the same underlying assets while varying their structure and issuance frequency, so changes in the menu reflect design choices alone. On the demand side, these studies observe flows at the product level; in our setting every contract and trade settles on a public blockchain, so wallet-level histories identify the traders who adopt each design, the designs from which they arrive, and the returns they subsequently earn. We can therefore follow a platform’s contract choices through to the outcomes of the traders who adopt them, and we find that the designs that best serve the platform’s volume objective are also those under which trader wealth erodes most rapidly.

Second, we contribute to the literature on prediction markets. Classical work asks whether prices aggregate information and whether they can be interpreted as probabilities ([Wolfers and Zitzewitz, 2004](#); [Manski, 2006](#); [Snowberg and Wolfers, 2010](#)), and recent evidence finds that modern platforms forecast macroeconomic releases and corporate earnings about as well as, or better than, surveys and professional analysts ([Diercks, Katz, and Wright, 2026](#); [Gomez Cram et al., 2025](#); [Rabetti, Shao, and Zhang, 2026](#)). A second group of papers studies who wins and who loses. On Polymarket, most wallets end with a loss and the top percentile of profitable wallets collects about three quarters of aggregate gains ([Akey et al., 2026](#); [Reichenbach and Walther, 2025](#); [Oliphant, 2026](#)); the winners are a small set of accounts that provide liquidity and execute better ([Akey et al., 2026](#); [Della Vedova, 2026b](#)) or, on another reading of the same ledger, an informed minority whose profits the crowd funds ([Gomez Cram et al., 2026](#)). On Kalshi, which charges fees, losses concentrate among takers of liquidity and in cheap longshot contracts ([Bürigi, Deng, and Whelan, 2025](#)), and market makers earn least on contracts written on asset prices, indices, and macroeconomic releases, where private information is dispersed ([Bartlett and O’Hara, 2026](#)). A third group studies the integrity of listed markets: insider trading

(Mitts and Ofir, 2026; Liu, 2026; Della Vedova, 2026a; Nechepurenko, 2026), settlement manipulation (Dai, Jia, and Yu, 2026; Abid, 2026), wash trading (Siroly et al., 2025), and the microstructure of blockchain-settled markets (Tsang and Yang, 2026; Dubach, 2026). Throughout this work, the menu of tradable contracts is taken as given. We study instead how the menu itself is chosen. The choice matters for every question above: Galanis, Ioannou, and Kotronis (2024) show theoretically that the design of the traded security determines whether a prediction market aggregates information at all, and Dai, Jia, and Yu (2026) show that five-minute contracts—the design in which we find trading eventually concentrates—allow manipulation of the settlement price, and recommend a longer horizon as the remedy. Both papers evaluate given designs; we explain why a volume-maximizing platform chooses them, and document that Polymarket’s menu evolved as the theory predicts.

Third, we contribute to the literature on how the trading environment shapes retail investors’ behavior and outcomes. Heimer and Simsek (2019) show that capping retail leverage reduces speculative losses without impairing liquidity; Barber et al. (2022) show that brokerage app design steers attention and order flow; Bryzgalova, Pavlova, and Sikorskaya (2023) document that retail options traders concentrate in the cheapest, shortest-dated contracts and lose money on average; Chapkovski, Khapko, and Zoican (2026) show that gamified interfaces raise trading intensity. The demand side of this environment is well documented: retail investors favor lottery-like payoffs (Barberis and Huang, 2008; Kumar, 2009; Boyer and Vorkink, 2014; Agarwal et al., 2025; Beckmeyer, Branger, and Gayda, 2023), and their performance improves slowly, if at all, with experience (Odean, 1999; Barber et al., 2020; Chague, De-Losso, and Giovannetti, 2019). In these studies, the trading environment varies while the set of tradable contracts is fixed. We study the margin that precedes all of them: which contracts exist to be traded. Over our sample, Polymarket’s menu is the platform’s own choice, and it moves from long-dated structured claims to five-minute binary bets—a compression parallel to the migration of retail options trading toward weekly and zero-days-to-expiry contracts (Beckmeyer, Branger, and Gayda, 2023). We show that this expansion redirects trading, raises its intensity, and lowers its returns: contract supply is a distinct instrument through which platforms shape retail outcomes.

3 Setting, Data, and Contract Taxonomy

3.1 Setting

Polymarket is a decentralized prediction market that operates on a public blockchain and settles in a dollar-pegged stablecoin, USDC. Every contract, order, and transfer is recorded on-chain and can be reconstructed without relying on proprietary exchange records.

The tradable unit is a *market*: a binary contract that pays one dollar if a stated proposition resolves true and zero otherwise, so that its price is directly interpretable as a probability. Markets are organized in a three-level hierarchy. Related markets are bundled into an *event*, which is the object a user actually browses and comments on; an event may, for instance, collect the full ladder of strike-price markets written on the same underlying for the same expiry. Events that recur on a fixed template are in turn grouped into a *series*, which is the exchange’s device for issuing the same contract design over and over—a new one every five minutes, every hour, or every week.

The series layer is central to our analysis because it makes the platform’s contract creation directly observable: each series records every instance in which the platform relists the same contract design. Series names typically encode both the object of prediction and the frequency with which new contracts are issued, a convention we exploit below to classify the market universe.

New contracts are not introduced automatically; instead, they follow a centralized listing process. Although users may submit proposals for new contracts, only a subset is ultimately approved for trading.² Each proposal must specify a clearly defined market question, an objective resolution source, and supporting evidence of investor demand. The exchange then determines whether and when to list the proposed contract, taking into account expected trading activity and other commercial considerations. As a result, the supply of tradable contracts is an endogenous outcome of the platform’s listing process rather than a mechanical reflection of real-world events.

²Recent market proposals submitted to Polymarket on X are often tied to discrete outside events—for example, whether Lil Durk will be found guilty on at least one count in his federal trial, the sex of Charles Leclerc and Alexandra Saint Mleux’s first child, or next-point outcomes during the US Open. The financial contracts we study differ in a fundamental respect. Their underlying assets trade continuously, while the platform chooses how to package them into recurring contracts—including the progressive introduction of Up/Down markets at shorter horizons. Menu expansion in our financial sample therefore primarily reflects the platform’s choice of how to create repeated trading opportunities on an existing underlying, rather than the arrival of new events to trade.

3.2 Data

We assemble our dataset directly from Polymarket’s public Application Programming Interfaces (APIs). Specifically, we use the Gamma API to collect contract metadata and the Central Limit Order Book (CLOB) API to obtain transaction-level trading records from January 2023 through March 2026. Although the ledger begins in January 2023, the financial series that define our sample become active only in February 2024, so the analysis effectively covers February 2024 through March 2026. We first retrieve the complete universe of markets, events, and series, together with their descriptive information, including category tags, creation and resolution timestamps, resolution status, and trading volume. We then collect the complete transaction history of every wallet that has ever traded on the platform. The resulting trade ledger contains approximately 1.04 billion transaction records, each reporting the traded market and outcome, the trade direction, quantity, and execution price, and a timestamp. We further retrieve all open and settled positions for every wallet; these records provide an independent consistency check on the reconstructed trading histories and supply settlement prices for all resolved contracts.

Our raw data include both unresolved contracts and a small number of positions with incomplete trading histories. To ensure the accuracy of realized payoff calculations, we apply two sample restrictions. First, we retain only resolved contracts, defined as markets with valid resolution outcomes and timestamps. This restriction is required for all analyses involving realized trading performance and retains approximately 93% of the market universe. Second, we exclude positions with incomplete or inconsistent trading records. In particular, we remove positions whose reconstructed trading histories cannot be reconciled with the exchange’s position records, as well as those for which the recorded sales exceed the shares previously purchased.

We further restrict the analysis to series-based markets linked to underlying financial assets. Compared to one-off event markets, these markets follow highly standardized contract templates. Because the set of underlying assets is essentially fixed, the platform expands the investment opportunity set along contract terms alone—maturity, strike thresholds, and issuance frequency—so that variation in the menu reflects design choices rather than changes in the underlying source of uncertainty. In addition, the settlement value of a financial contract is a public price, so private information about

the outcome plays a smaller role than in event contracts: on Kalshi, contracts on asset prices, indices, and macroeconomic releases show among the lowest informed price impact of any category (Bartlett and O’Hara, 2026), and the insider-trading cases documented on Polymarket concern political, geopolitical, and corporate events (Mitts and Ofir, 2026; Liu, 2026). Trading in financial contracts is thus closer to the heterogeneous-beliefs setting of our model, and this recurring issuance of near-identical contracts provides a clean setting for studying the effects of menu expansion. This restriction defines the scope of our claims: we characterize the design of the platform’s menu of financial contracts and its consequences, not the platform in its entirety. Political and sports markets are excluded not because they are small—they carry much of the platform’s volume—but because their listing follows the arrival of outside events, so their menu reveals little about design choices.

We classify markets into cryptocurrency, stock price, earnings, index, commodity, and foreign-exchange categories using the descriptive tags attached by Polymarket to each event. These tags are assigned by the exchange itself and largely follow the category structure used to organize markets on the platform. We adopt the platform’s taxonomy as the basis for our classification, while making a small number of adjustments where temporary or promotional categories create inconsistencies in the underlying hierarchy. Tags are available for 98.7% of all events, providing nearly complete coverage of the market universe.

Our classification proceeds in four steps. First, we remove administrative and frequency tags that do not describe the underlying subject of a contract. Second, we map each informative tag to a predefined category and, where applicable, a subcategory. The mapping is based primarily on the platform’s taxonomy, but we reassign some overly broad labels such as business or finance to more specific asset-class groupings like stocks, earnings, indices, commodities, or foreign exchange when the event title and slug make the underlying subject clear.

Third, when an event carries tags from multiple categories, we assign it to a single category according to a fixed priority ordering—Crypto, Sports, Weather, Economy & Finance, Geopolitics, Politics, Tech, Culture—so that, for example, an event tagged as both Crypto and Politics is classified as Crypto. Finally, for the small fraction of events without informative tags, we assign categories using keywords extracted from the market title and identifier. Events that cannot be classified through either

procedure are excluded from the sample.

3.3 Contract Structure and Issuance Frequency

Series-based markets—the overwhelming majority of both listed markets and trading volume on the platform (see Appendix Figure A.1)—are organized under standardized naming conventions that encode key contract characteristics. We exploit these conventions to classify each market along three dimensions—contract structure, issuance frequency, and underlying asset—which together characterize how the platform expands the investment opportunity set through the repeated introduction of contracts of the same design. The first dimension, contract structure, describes the type of prediction the contract offers; we describe the five structures in the order in which the platform first introduced them. *Thematic* contracts, the platform’s earliest designs, ask individually specified questions about an underlying asset. Because they fit none of the standardized templates described below, we group them in a residual category, labeled *Other* in the tables. *Hit Price* contracts ask whether the underlying asset will reach a specified price threshold at any point before expiration. The platform subsequently created richer contract structures. *Multi-strike* contracts simultaneously list multiple ordered price thresholds for the same underlying and maturity, allowing investors to trade different price levels independently. For example, contracts asking whether Bitcoin will exceed \$100,000, \$110,000, or \$120,000 by the same expiration date may all resolve positively because the thresholds are nested. By contrast, *Neg-risk (multi-outcome)* contracts partition the possible realizations of a single prediction into mutually exclusive outcomes. For example, rather than asking whether Bitcoin will exceed a sequence of price thresholds, the platform may ask which of several non-overlapping price intervals Bitcoin will finish in at expiration. Since these intervals are mutually exclusive, exactly one contract resolves positively. *Up/Down* contracts, the last structure the platform introduced, ask whether the underlying asset will rise or fall over a specified horizon. Unlike hit-price contracts, the prediction concerns the terminal direction of the price rather than whether a level is ever attained, and each contract contains a single binary prediction—the simplest structure on the menu.

The second dimension, *issuance frequency*, describes how often the platform reissues contracts based on the same prediction template. Frequency is a distinct margin of expansion: without intro-

ducing a new underlying asset or a new payoff structure, the platform can substantially enlarge the set of tradable contracts simply by shortening the interval between successive launches. We classify series as monthly, weekly, daily, 4-hour, hourly, 15-minute, or 5-minute—the order in which the platform introduced these frequencies—with a residual category for event-driven contracts that are not launched on a regular schedule.

Finally, we assign a canonical underlying asset to each series. For cryptocurrency contracts, ticker abbreviations and alternative names are mapped to a canonical asset identifier. For equity, earnings, index, commodity, and foreign-exchange contracts, the reported ticker serves as the canonical identifier. This allows us to distinguish the introduction of new underlying assets from the multiplication of contracts based on existing assets. The timing of these introductions—when each structure and each frequency entered the menu, and in which asset class—is itself a central object of the analysis, and Section 5 documents it in full.

4 A Model of Platform-Optimal Contract Design

This section develops a model of a trading platform that designs the securities its users trade. The platform cannot create the disagreement that drives trading, but it controls the features of the contract through which disagreement is expressed: what the payoff loads on, and how short its horizon is—and hence how often it is relisted. The trader side of the model follows [Simsek \(2013b\)](#), in which positions decompose into a risk-sharing component and a speculative component; the platform side follows the exchange-design tradition of [Duffie and Jackson \(1989\)](#), in which the venue selects the contract that maximizes trading volume. The analysis proceeds in four steps. Section 4.3 shows why a volume-maximizing platform loads its contracts on the disputed outcome, and Section 4.4 shows that the same motive compresses the horizon. Section 4.5 traces the consequences of the resulting contracts for trader wealth, and Section 4.6 derives the inflow condition on which the platform’s continued operation depends. Section 4.7 maps each result to its empirical counterpart, and Appendix C collects the derivations in full.

4.1 Environment

Following [Simsek \(2013b\)](#), consider two groups of traders, *optimists* and *pessimists*, indexed $g \in \{O, P\}$, each of unit mass, with mean–variance preferences and common risk-aversion parameter $\theta > 0$. Uncertainty has two orthogonal components. The first, v_1 , is a *disputed* outcome about which the groups hold different beliefs: optimists assign mean $+D$ and pessimists mean $-D$, so $D \geq 0$ measures the exogenous degree of disagreement, which the platform cannot alter. The second, v_2 , is an *agreed* source of risk on which both groups share the same (mean-zero) belief. The two components are independent, and we normalize their variances to unity:

$$\mathbb{E}_O[v_1] = +D, \quad \mathbb{E}_P[v_1] = -D, \quad \mathbb{E}_g[v_2] = 0, \quad \text{Var}(v_1) = \text{Var}(v_2) = 1. \quad (1)$$

Traders also carry opposite background exposures. A trader in group g is endowed with a background position:

$$b_g = s_g (v_1 + \alpha v_2), \quad s_O = +1, \quad s_P = -1, \quad (2)$$

where $\alpha \geq 0$ measures the importance of genuine hedging demand. We interpret the background exposure as holdings of the underlying itself—for instance, spot cryptocurrency positions—so that these contracts could in principle serve as insurance. When α is small, traders have little agreed risk to share; when α is large, the agreed component looms large.

Before trading begins, the platform chooses a contract with payoff

$$a = c_1 v_1 + c_2 v_2, \quad c_1^2 + c_2^2 = 1, \quad (3)$$

issued in zero net supply—as on Polymarket, where each market is a fully collateralized side bet in which one side’s payment is the other’s receipt. The design vector (c_1, c_2) fixes what the security pays off on, and the normalization fixes its scale, so design choices concern the direction of exposure rather than its magnitude. The closer c_1 is to one, the closer the contract is to a pure directional bet on the disputed outcome: the pure bet v_1 , with $(c_1, c_2) = (1, 0)$, is what a prediction market lists, while a contract aligned with the background exposure, $(c_1, c_2) \propto (1, \alpha)$, is a compound claim that primarily

shares risk.

4.2 Equilibrium Trading

A trader's net worth is cash plus the payoff on the background position plus the payoff on the contract position x_g taken at price p ,

$$n_g = e_g + b_g + x_g (a - p). \quad (4)$$

Each trader chooses x_g to maximize a mean–variance objective under her own beliefs,

$$\max_{x_g} \mathbb{E}_g[n_g] - \frac{\theta}{2} \text{Var}_g[n_g]. \quad (5)$$

Solving (5) and imposing market clearing (Appendix C) sets the price equal to the component of the payoff on which the two groups agree, $p = 0$, and yields equilibrium positions that decompose into a hedging and a speculative term,

$$x_O = \underbrace{-(c_1 + \alpha c_2)}_{\text{hedging}} + \underbrace{\frac{c_1 D}{\theta}}_{\text{speculation}}, \quad x_P = -x_O. \quad (6)$$

The first term is hedging demand: it would be present even if the two groups shared identical beliefs ($D = 0$), and it offsets the background exposure (2). The speculative position

$$x_S = \frac{c_1 D}{\theta} \quad (7)$$

is the central object in what follows: it rises with disagreement D , with risk tolerance $1/\theta$, and with the contract's weight on the disputed outcome c_1 —a more speculative contract makes each trader bet more.

Two features of this equilibrium are worth highlighting. First, payoff components that traders agree on are priced rather than traded: in a binary market the equilibrium price is the consensus probability, so disagreement moves positions rather than prices. Second, because the two groups take exactly opposite positions, the speculative bet is purely redistributive—aggregate trader wealth is unchanged within any single contract, and larger disagreement simply produces larger opposing bets.

4.3 The Platform's Optimal Contract

The platform's objective throughout is current trading volume, the metric on which platforms of this kind are valued; Section 4.6 derives the consequences of this objective for long-run viability. Since $x_P = -x_O$, total volume is proportional to $|x_O|$, and, holding the participant base fixed, the platform favors the design direction (c_1, c_2) on the unit circle that maximizes the size of the equilibrium position,

$$\max_{c_1^2 + c_2^2 = 1} |x_O| = \left| c_1 \left(\frac{D}{\theta} - 1 \right) - \alpha c_2 \right|. \quad (8)$$

The maximand is the absolute value of a linear form on the unit circle, so it is maximized when (c_1, c_2) points along the coefficient vector (Appendix C). The platform's optimal contract is therefore

$$A^*(D) \propto \left(\frac{D}{\theta} - 1 \right) v_1 - \alpha v_2. \quad (9)$$

Proposition 1 *The optimal loading on the disputed outcome, $c_1^* = (D/\theta - 1)/\sqrt{(D/\theta - 1)^2 + \alpha^2}$, is strictly increasing in disagreement D and in risk tolerance $1/\theta$: it rises from $-1/\sqrt{1 + \alpha^2}$ at $D = 0$, through zero at $D/\theta = 1$, to one as $D/\theta \rightarrow \infty$. Greater disagreement thus tilts the contract toward the pure disputed bet v_1 : catering to disagreement is equivalent to designing contracts that emphasize disputed outcomes.*

Equation (9) is most transparent at the two extremes of disagreement. With no disagreement ($D = 0$) the optimal contract is proportional to $v_1 + \alpha v_2$, a replica of the background exposure: the volume-maximizing product is a hedging instrument. As D/θ grows, the hedging ingredient shrinks relative to the disputed one, and in the limit the contract converges to a bare directional bet. The platform's menu spans the range between these two extremes. Designs whose payoffs are graded in the size of the price move—multi-strike ladders and hit-price contracts—carry more of the agreed, level-dependent component and come closest to an insurance product: a position in them can offset a holder's exposure to the underlying. At the opposite extreme sits the at-the-money binary directional contract, particularly at high issuance frequencies: it pays the same dollar whether the move is one basis point or five percent, and its horizon is far shorter than any exposure a hedger might wish to

offset.³

This mechanism is particularly relevant for prediction markets because they combine two distinctive features. First, prediction markets are explicitly designed for trading on differences in beliefs about future events, making disagreement an intrinsic source of trading activity rather than a by-product of security issuance. Second, contract creation is endogenous: the platform continuously determines which uncertainties become tradable and how they are packaged into securities. Together, these features make Polymarket a natural setting for understanding how platform incentives interact with investor disagreement to shape contract creation and, ultimately, the organization of market activity.

4.4 Faster Contracts: The Horizon Margin

Proposition 1 concerns which dimension the contract spans. The platform also controls the contract’s *horizon*, and the same volume motive that loads the payoff on the disputed outcome also compresses the horizon. Let contracts have horizon $h > 0$, and normalize calendar time to unit length, so that a horizon- h contract can be relisted $1/h$ times, with disagreement realized and renewed each cycle. Because a shorter horizon resolves a smaller slice of the disputed outcome, per-cycle disagreement is $D(h)$ with

$$D(0) = 0, \quad D'(h) > 0, \quad D''(h) < 0. \quad (10)$$

For binary contracts these properties can be derived rather than assumed. Following the heterogeneous-beliefs framework in which traders disagree about expected returns but agree on volatility (Basak, 2005; Hong and Stein, 2007), suppose the two groups disagree about the drift of the underlying price, expecting $+\mu$ and $-\mu$ per unit of time, while agreeing on its volatility σ . Over a horizon h the price change is approximately normal with mean $\pm\mu h$ and standard deviation $\sigma\sqrt{h}$, so the probability that group g assigns to the price ending up is $\Phi(\pm(\mu/\sigma)\sqrt{h})$, and the disagreement about the binary outcome is

$$D(h) = \Phi\left(\frac{\mu}{\sigma}\sqrt{h}\right) - \Phi\left(-\frac{\mu}{\sigma}\sqrt{h}\right) \approx \sqrt{\frac{2}{\pi}} \frac{\mu}{\sigma} \sqrt{h} \quad \text{for small } h, \quad (11)$$

³The model attributes demand entirely to belief disagreement and abstracts from preferences over payoff shapes. Retail investors appear to value simple binary bets for their own sake: Goodman and Puri (2025) find that option traders overpay for all-or-nothing payoffs relative to spreads that dominate them state by state. Such a preference would reinforce the platform’s tilt toward Up/Down contracts over multi-strike ladders; we leave it out to show that disagreement alone generates the platform’s design incentives.

which is zero at $h = 0$, increasing, concave, and bounded by one. The square root reflects that short-horizon price moves are dominated by volatility, which scales with $\sqrt{h}$, rather than by drift, which scales with h : a five-minute contract contains almost nothing to disagree about, which is precisely why it must be reissued so often.

From the speculative term in (6), single-cycle volume is proportional to $c_1 D(h)/\theta$; aggregating over the $1/h$ cycles per unit time, the platform's volume flow is

$$V(h) = \frac{1}{h} \cdot \frac{c_1 D(h)}{\theta} \propto \frac{D(h)}{h}. \quad (12)$$

Since D is concave through the origin, the average product $D(h)/h$ strictly exceeds the marginal product $D'(h)$ and is decreasing in h , so $V(h)$ is maximized as $h \rightarrow 0$: faster contracts renew disagreement, and the trading it generates, more often, and this renewal effect dominates the thinner per-cycle disagreement. In practice the platform incurs a fixed cost $F > 0$ per listing—oracle specification, liquidity seeding, attention. Its revenue, in turn, is proportional to trading volume: let $\tau \in (0, 1)$ denote the proportional trading cost, the fraction of each dollar of volume that accrues to the platform and its liquidity providers through explicit transaction fees and effective bid–ask spreads. From the traders' side τ is a transaction cost; from the platform's side it is revenue per unit of volume. Writing the platform's per-unit-time payoff as revenue net of listing cost, with $m \equiv \tau c_1/\theta$,

$$\Pi(h) = \frac{1}{h} [m D(h) - F], \quad (13)$$

the first-order condition gives a unique interior optimum h^* satisfying

$$m [D(h^*) - h^* D'(h^*)] = F. \quad (14)$$

Proposition 2 (Contraction of the horizon) *The volume-maximizing horizon h^* defined by (14) is decreasing in the trading cost τ and in the loading c_1 , and increasing in the listing cost F . In the frictionless limit $F \rightarrow 0$, $h^* \rightarrow 0$.*

Proposition 2 ties the platform's trajectory to the cost of listing: as the technology for creating

markets is standardized and automated, F falls and the optimal horizon converges to 0. Under the binary specification (11), h^* is proportional to F^2 , so even moderate declines in listing costs produce sharp compression, and underlying assets that generate more drift disagreement are relisted faster. A platform optimizing on both margins therefore converges on a menu of short-dated directional bets on the disputed outcome.

4.5 The Distribution of Trader Profits

The equilibrium of Section 4.2 is purely redistributive: within any single contract, one side's gain is the other's loss, and it is tempting to infer that traders as a group break even. Once contracts are repeated—which, by Proposition 2, is exactly what short horizons deliver—this inference fails, and in a specific direction: the typical trader loses while a shrinking minority accumulates the winnings, even when trading itself is costless. Consider a trader who stakes a fraction $f \in (0, 1)$ of current wealth on the disputed outcome each cycle, so that her wealth is multiplied by $(1 + f)$ or $(1 - f)$ with equal probability. Full collateralization makes this the natural description of behavior on the platform we study; we treat it as a positive staking rule rather than an implication of the preferences in (5). The bet is fair in *arithmetic* terms, $\mathbb{E}[W_{t+1}/W_t] = \frac{1}{2}(1 + f) + \frac{1}{2}(1 - f) = 1$. But wealth compounds multiplicatively, and the long-run fate of a trader is governed by the geometric growth rate,

$$g \equiv \mathbb{E}[\ln(W_{t+1}/W_t)] = \frac{1}{2} \ln(1 + f) + \frac{1}{2} \ln(1 - f) = \frac{1}{2} \ln(1 - f^2) < 0. \quad (15)$$

The gap between the arithmetic mean (zero) and the geometric mean (strictly negative) is volatility drag, a consequence of the concavity of the logarithm and present for any $f > 0$. By the strong law of large numbers, $t^{-1} \ln(W_t/W_0) \rightarrow g < 0$ almost surely, so $W_t \rightarrow 0$ for almost every trader, while aggregate wealth is conserved each period; the reconciliation is that a vanishing fraction of traders—those on long lucky streaks—come to hold a growing share of the total.

The horizon of Section 4.4 governs how often this drift is incurred. Under the normalization there, a horizon- h design completes $1/h$ cycles per unit of calendar time, so by calendar date T a trader has been through $t = T/h$ rounds of redistribution and her expected log wealth is $(T/h)g$. The stake f sets the drift per cycle; the horizon sets the number of cycles the calendar delivers, and hence the

speed at which the drift accumulates in real time. In particular, the median trader’s wealth at date T is $W_0 e^{gT/h}$ exactly, falling by half every $h \ln 2/|g|$ units of calendar time.

Proposition 3 (Erosion of trader wealth) *In the repeated symmetric market with proportional stake $f > 0$, horizon $h > 0$, and no trading cost, the fraction of traders whose cumulative wealth falls below their starting wealth converges to one as calendar time $T \rightarrow \infty$. The wealth distribution at date T depends on the horizon only through the number of completed cycles T/h : the calendar time required for any given erosion of trader wealth is proportional to h . A per-cycle trading cost $\tau > 0$ shifts the drift to $g - \tau < g$, accelerating the loss at any horizon.*

Two points deserve emphasis. First, the result in Proposition 3 requires no difference in skill or information across traders: it is a pure consequence of repeated multiplicative redistribution, and skill dispersion sharpens the tail without being needed for the result. Second, the mechanism compounds Proposition 2: the drift g does not depend on h , but the median trader’s halving time $h \ln 2/|g|$ does, so the same staking behavior that erodes wealth over years on monthly contracts does so within hours on five-minute ones. The model thus predicts that as the menu shifts toward faster designs, the share of accounts in profit should fall and the time it takes a fixed amount of trading capital to be exhausted should shrink with the contract’s horizon—the objects measured in Section 5.4.

4.6 A Capital-Flow Condition

Proposition 3 has a consequence for the platform itself: the contracts that maximize current trading volume also consume the wealth that traders can continue to stake. Shorter horizons compress more trading cycles into a given amount of calendar time, raising current volume while shortening the time over which a given stock of trader capital can support it, so a platform that optimizes current volume may progressively deplete its own trading base. The platform can sustain volume only if new capital replaces what is lost. Let K_t denote aggregate live trading capital at the start of period t , and let $d \in (0, 1)$ denote the fraction of that capital that leaves the pool each period: the trading cost τ diverts part of every dollar of volume to the platform and its liquidity providers, and the winning minority of Proposition 3 withdraws part of its realized gains. Both outflows move with trading volume, and trading volume moves with the capital deployed, so a faster menu—which by (12) generates more

volume from the same capital—corresponds to a larger d . With gross inflow $I_t \geq 0$ from newly funded and returning accounts, the pool evolves as

$$K_{t+1} = K_t (1 - d) + I_t. \quad (16)$$

Observation 1 (Capital-flow condition) *The capital pool, and with it volume and platform revenue, is non-decreasing period over period if and only if inflow covers the outflow, $I_t \geq d K_t$. If, from some date on, inflow falls short of this threshold and decays geometrically, then capital, volume, and platform revenue decay geometrically with it.*

The requirement does not vanish at $\tau = 0$: d remains strictly positive because winners withdraw gains from the pool regardless of the trading cost; the binding constraint on a low-fee platform is therefore the pace of new-account acquisition—a discipline directly relevant to our setting, since Polymarket charged no explicit fees for most of the sample period, introducing them first on the cryptocurrency designs near the end of our sample window. The horizon determines how demanding the requirement is: because a faster menu raises d , the inflow needed to sustain volume, $d K_t$, grows as the menu compresses toward shorter horizons.

4.7 Discussion: Mapping the Model to the Evidence

Taken together, the three propositions and the capital-flow condition describe the platform’s trajectory. The volume motive loads the contract on the disputed outcome (Proposition 1) and compresses the horizon (Proposition 2); short horizons multiply the number of redistributive cycles per unit of calendar time, and repeated redistribution erodes the typical trader’s wealth (Proposition 3); the winners’ withdrawals, together with the trading cost τ , form precisely the outflow of capital that new inflow must replace (Observation 1). The upshot is a tension internal to the platform’s own optimization: faster and more speculative contracts raise near-term volume and revenue, but they do so by accelerating the erosion of trader wealth, and hence the inflow of new capital that sustaining that volume requires. The capital-flow condition states what a forward-looking platform would weigh against current volume: the inflow that its menu must attract to sustain that volume.

The mapping to the empirical analysis is then direct. The menu’s convergence on fast directional designs tests Propositions 1 and 2. Along the structures, designs whose payoffs are graded in the size of the price move—multi-strike ladders and hit-price contracts—tie part of the payoff to the magnitude and shape of the return distribution, on which traders broadly agree, while Up/Down contracts condition on direction alone; to the extent that disagreement concentrates in the direction of the move rather than its magnitude, the migration toward Up/Down raises the loading c_1 on the disputed component—the movement Proposition 1 predicts when disagreement is large relative to the hedging motive. Along the frequencies, the move from monthly to five-minute issuance is the compression of the horizon h that Proposition 2 predicts as the cost of listing falls. The per-wallet performance measures—mean returns, the share of wallets ending in profit, and the time for a fixed amount of trading capital to be exhausted—should worsen for more speculative and faster designs (Proposition 3; Section 5.4). Finally, the recruitment of new accounts on which the capital-flow condition makes continued volume depend is the subject of Section 5.5.

5 Empirical Analysis

We bring the model to the Polymarket data in three steps. We first document how the contract menu evolved and summarize the market universe it has produced; we then trace how trading activity migrated across designs as they appeared; and we finally measure the consequences of the designs for the traders—their returns, their survival, and their replacement.

5.1 Contract Creation

Payoff components that traders agree on are priced rather than traded, so a platform whose revenue scales with volume should progressively strip them out and concentrate contract payoffs on the outcomes traders disagree about (Proposition 1). The simplest such contract is a binary bet on the direction of a price move, and the shorter its horizon, the more often the disagreement—and the trading it generates—renews (Proposition 2). A platform that adjusts its menu over time should therefore converge on this design, at the pace at which its listing costs fall: the longer it has been optimizing an asset class’s menu, the simpler and faster the contracts it lists. Our sample is built to make this

prediction testable: because it restricts attention to financial underlyings, whose menu changes reflect platform design choices rather than the arrival of new events (Section 3.2), the sequence in which structures and frequencies are introduced traces the platform’s choices directly. We therefore begin by characterizing how Polymarket expands its menu of financial contracts, providing, to our knowledge, the first empirical account of a platform’s sequence of contract-design choices.

Panel (a) of Figure 1 records when each contract structure was first introduced, separately for each asset class. For each structure, a marker locates the date the platform first listed it within an asset class, and the connector spans the lag between the earliest and latest class to adopt it. The figure reveals substantial differences in the timing with which contract structures are introduced across asset classes. In cryptocurrency, the platform introduces them in a clear sequence: thematic contracts first (February 2024), followed by hit-price (October 2024), multi-strike (December 2024), neg-risk (February 2025), and finally up-or-down contracts (March 2025). Equities and commodities follow the same broad ordering but with a substantial lag. Thematic contracts again appear first—in September 2024 for equities and February 2025 for commodities—while the remaining structures arrive only in late 2025; in equities, up-or-down contracts are introduced somewhat earlier than the hit-price, multi-strike, and neg-risk structures, which cluster in October and November 2025. The shaded band marks this late-2025 period, when stock-related and commodity contracts rapidly broaden their menu of payoff structures. Overall, the platform expands the payoff dimension one structure at a time rather than all at once, and the cryptocurrency sequence leads the other asset classes by roughly a year.

Panel (b) of Figure 1 complements panel (a) by tracing the evolution of contract frequency. The y-axis orders issuance frequencies from one-off contracts to five-minute contracts and records when each asset class first reaches each frequency. Cryptocurrency moves rapidly through this sequence, beginning with one-off and monthly contracts before successively introducing weekly, daily, four-hourly, hourly, fifteen-minute, and finally five-minute markets. In calendar time this compression is fast: crypto’s first one-off contracts date to February 2024, with monthly series following in October 2024, weekly in December 2024, daily in March 2025, four-hour and hourly in May 2025, fifteen-minute in September 2025, and five-minute contracts in February 2026—a horizon so short that no scheduled disclosure or fundamental information bears on the outcome, which is why [Della Vedova \(2026a\)](#) uses

such markets as a placebo class in tests for informed trading, and little of the underlying’s own price discovery occurs before settlement (Dai, Jia, and Yu, 2026).⁴⁵ By contrast, equities and commodities follow the same broad sequence but with a substantial delay—the first stock and earnings series appear only in September 2024—and have so far stopped at the daily frequency. Thus, while the direction of frequency expansion is broadly similar across asset classes, its intensity is overwhelmingly driven by cryptocurrency markets.

Panel (c) of Figure 1 crosses payoff structure with issuance frequency, with marker size proportional to the number of markets in each cell. The plot shows that the two design margins are not interchangeable: how short a contract’s issuance interval can become is determined by its payoff structure. Weekly is the only frequency at which all five payoff structures coexist. Past that point the richer designs progressively drop away: hit-price contracts stop at the daily frequency, neg-risk contracts appear only at the weekly and four-hour frequencies, and multi-strike contracts reach no further than the hourly frequency. Only the binary up-or-down design spans all eight frequencies, and at the fifteen- and five-minute horizons it is the only structure on offer—precisely the cells with the largest markers in the entire grid. The distribution of mass tells the same story: the multi-strike and neg-risk designs have their largest cells at the weekly-to-hourly frequencies, while the up-or-down row is thin at slow horizons and swells as the issuance interval shrinks. Equities and commodities, which have not yet progressed beyond the daily frequency, fill only 10 and 5 of the 40 possible structure-frequency combinations, against 21 for cryptocurrency. Frequency expansion, in short, is not a general property of the contract menu but a property of its simplest design: the platform reaches ever-shorter horizons only for simpler, binary bets.

The ordering mirrors the loadings of Section 4.3: payoffs graded in the size of the price move carry more of the agreed component, while the up-or-down bet conditions on direction alone, so, insofar as disagreement concerns direction, descending the structure sequence concentrates the payoff on the

⁴Two crypto series slugs, btc-multi-strikes-weekly and ethereum-multi-strikes-weekly, contain 53 markets dating back to October 2020. These are individually named, non-recurring contracts that were retrospectively assigned to slugs which only became recurring weekly ladders in late 2024. Before then, only 4% of contracts lasted about one week; afterward, 98% did. Retaining these markets would therefore misdate the weekly frequency to 2020, so we exclude them rather than relabel them.

⁵We exclude the 45,719 five-minute markets created before 12 February 2026 because they were effectively inactive, recording just 148 trades and \$8,486 in total volume. Trading began only on 12 February 2026, when activity jumped to 1.3 million trades in a single day. We therefore treat 12 February 2026 as the introduction of five-minute markets throughout the paper. All other designs became active immediately after listing.

disputed outcome (Proposition 1). Proposition 2 makes the interaction with frequency explicit: the optimal horizon falls with the loading c_1 , so the design most concentrated on the disputed outcome is precisely the one the platform relists fastest.

Panel (d) plots every market by listing date and issuance frequency. The black line traces the monthly median frequency of new listings: for each month, we rank that month’s newly listed markets by issuance frequency, from one-off to five-minute, and record the frequency of the median market. The scatter conveys two features of the expansion that a summary series alone would miss. First, new frequencies add to the menu rather than replace it: once a frequency appears, listings at that horizon continue almost without interruption through the end of the sample, so the menu accumulates frequencies as the platform moves toward shorter issuance intervals. Second, the mass of new listings nonetheless shifts steadily toward shorter horizons. Over time the cloud thickens at daily and sub-daily frequencies, and the median declines almost monotonically: newly listed markets are typically one-off contracts through September 2024, monthly from October 2024, weekly from March 2025, daily in May 2025, hourly from June 2025, fifteen-minute from November 2025, and five-minute by March 2026. In little more than two years, the representative new listing thus moves from a bespoke contract with no fixed expiration cycle to one that expires within five minutes.

The only interruption occurs in September 2025, when a burst of four-hour listings briefly raises the median before the downward trend resumes. The coloring of the scatter ties this migration back to panel (c): the weekly and daily bands mix several payoff structures, whereas the bands from the hourly frequency down are dominated by up-or-down contracts, and the fifteen- and five-minute bands contain no other structure at all. By the end of the sample the fifteen- and five-minute bands are the densest regions of the plot—the two fastest frequencies alone account for close to two thirds of all crypto markets ever listed—and high-frequency contracts have become the primary form of new market creation.

Taken together, the four panels of Figure 1 document precisely this convergence. Cryptocurrency takes it furthest: the menu expands one design at a time and, at each step, moves closer to the contract that the framework in Section 4 associates with maximal trading—a simple binary bet on a disputed outcome, renewed at ever-shorter horizons. Equities and commodities follow the same

sequence with a lag of roughly a year and—as of the end of our sample in March 2026—have stopped at the daily frequency, consistent with asset classes still in the process of converging to the same final menu. The direction of this innovation differs from that documented for retail structured products and specialized funds (Célérier and Vallée, 2017; Vokata, 2021; Ben-David et al., 2023). Those issuers earn a margin on each dollar sold, so they benefit from complexity that shrouds the markup or from themes that attract inflows. A platform whose revenue scales with trading activity sells no product and takes no markup; it benefits from contracts that repeatedly bring traders with opposing beliefs to the market. The relevant design margin is therefore not complexity but exposure to disagreement and the frequency with which it is renewed.

5.2 The Platform in Cross-Section

A new contract matters only if traders adopt it. Before turning to trading activity, we describe the market universe that the expansion has produced at three levels of aggregation: the series through which the platform organizes issuance, the markets that these series generate, and the wallets that trade them—the three objects whose interaction the model of Section 4 describes. Table 1 describes the first of these—how markets are organized into series—while Table 2 reports summary statistics for the resulting markets and for the wallets trading them.

We restrict attention to series-covered markets and classify each one by contract structure and by issuance frequency, separately for Crypto (Panel A), Stocks/Earnings/Indices (Panel B), and Commodities/Forex (Panel C). For each category, we report the number of distinct series, markets, and underlying tickers, the average number of traders per market, total trading volume, the average number of markets per ticker, and the average created-to-closed lifespan.

The table’s central message is that the three principal segments sit at different stages of the same expansion: cryptocurrency, where issuance has reached the fifteen- and five-minute frequencies, multiplies contracts written on a small set of underlying assets; equities, where issuance has not yet extended past the daily frequency, still grow mainly by adding underlyings; and commodities, where one-off contracts still carry nearly all volume, remain at an earlier stage still. In crypto, the exchange writes an enormous number of short-dated, repeating contracts on a small set of underlyings: 210,469

markets rest on just 124 series and 27 tickers, or nearly 8,000 markets per underlying asset. In equities, expansion instead proceeds by adding names rather than contracts: 9,881 markets are spread across 191 tickers—roughly 52 markets each—and the segment’s 173 Up/Down series cover 164 distinct underlyings, essentially one series per name. Most of this breadth comes from earnings rather than prices: 142 of the 173 series are written on quarterly earnings announcements, one per firm, while the price series concentrate on a few dozen large stocks and indices. Per underlying, contract issuance is thus about one hundred and fifty times as intense in crypto, and the money follows the contracts: cumulative volume reaches \$12.5 billion in crypto against \$205 million in equities, a ratio of about sixty to one.

Sub-panel A.1 shows that crypto issuance is concentrated in the simplest possible design. Up/Down contracts—single binary bets on whether a coin’s price will rise or fall—account for 168,774 markets, or four of every five crypto markets ever listed, with more than 14,000 contracts written on the average underlying. The older and richer designs play a different role: they list far fewer markets but draw larger crowds per contract, with hit-price and thematic markets averaging 554 and 661 traders each, against 408 for the average Up/Down market. The platform, in other words, can grow along two margins—designs that draw many traders to a few contracts, and designs that spread traders across an ever-expanding number of them—and its recent growth is overwhelmingly of the second kind.

Sub-panel A.2 traces the same expansion along the frequency margin, where the platform reissues the same contract designs at horizons ranging from monthly down to every five minutes. These fast frequencies dominate the count: fifteen-minute markets alone number 72,113—34% of every crypto market the platform has ever listed—with five-minute markets close behind at 61,920, or 29%. Together, the five- and fifteen-minute buckets contain 64% of all crypto markets, the fifteen-minute bucket is the platform’s single largest source of volume at \$3.8 billion, and contracts reissued at sub-daily frequencies capture roughly two thirds of crypto volume overall. The trade-off across frequencies mirrors the one across structures: the average monthly market attracts 1,593 traders and stays open for about a month, while the average five-minute market attracts 521 and disappears within a day. The lower participation per contract is more than offset by the number of contracts.

Panel B shows the equity segment at a much earlier stage of the same expansion. Issuance remains

anchored at low frequencies, with weekly and monthly contracts making up 73% of the segment’s markets, and one-off, event-driven contracts still carry much of the weight: 443 one-time markets—under 5% of the segment’s count—generate \$79 million, or two-fifths of its total volume, and draw by far its largest crowds, averaging 591 traders per market against 47 for weekly contracts. The equity segment at the end of our sample thus resembles the crypto segment before the platform began compressing its issuance intervals.

Panel C shows that expansion remains limited in commodity and foreign-exchange markets. The segment is small and recent, comprising 633 markets across 29 series and 17 underlying assets. Market counts and trading activity are concentrated in different contract types. Daily Up/Down contracts on commodity futures account for the largest number of markets (416) but generate little trading volume, whereas a small number of hit-price ladders on oil, gold, and silver account for almost all trading activity. This pattern is also reflected in issuance frequency. One-off contracts account for 208 markets, roughly a third of the segment, yet generate \$57 million of its \$60 million in trading volume and attract the largest participation, averaging 323 traders per market compared with 32 for the daily series. Short-horizon reissuance is thus present in this segment but has not taken hold, and new contracts are otherwise introduced only occasionally. As a result, neither the intensive reissuance observed in cryptocurrency markets nor the broad expansion across underlying assets characteristic of equity markets has yet emerged in this segment.

Table 2 reports cross-sectional summary statistics for the sample, with individual markets as the unit of observation in Panel A and individual wallets in Panel B. We first compute the value of each variable at the level of the individual market or wallet, and then report the distribution of these values across the 221,000 markets and 1.25 million wallets in our sample. For each variable, we report the number of observations used in the second step of the computations, the mean, standard deviation, and the 1st, 25th, 50th, 75th, and 99th percentiles.

Panel A documents two features of the market cross-section: trading activity is concentrated in a small fraction of markets, and the typical contract is short-lived. The median market attracts \$13,057 in volume, 946 trades, and 174 distinct traders over its life, but the means are far larger—\$57,843, 2,772, and 345, respectively—and the top percentile of markets clears more than \$565,000 in volume.

A small set of contracts thus accounts for a disproportionate share of platform activity.

Turning to longevity, the median market is open for just over one day, three quarters of markets survive no more than 2.1 days, and even the 99th percentile closes within a month. As documented in Table 1, this compressed life cycle is the direct imprint of the platform’s issuance model, which relists the same contract templates at frequencies as high as every five minutes. These cross-sectional averages, however, pool observations from the platform’s entire history, over which—as just documented—the composition of the contract menu changed almost entirely, so the figures in the table mix contract vintages with very different characteristics.

The remaining rows of Panel A characterize how these markets trade. Stakes are small: the median market’s average trade size is \$13, and in three quarters of markets the average trade moves about \$29. Trading is also essentially costless for most of the menu: the median market generates no measurable fee revenue at all and the average fee rate rounds to zero, so the frictions that might otherwise discourage rapid, repeated betting are largely absent. The zeros, however, are not uniform: the platform introduced taker fees on fifteen-minute cryptocurrency contracts in January 2026, applied them to five-minute contracts from their launch, and extended them to the remaining cryptocurrency designs in March 2026 (Dai, Jia, and Yu, 2026; Akey et al., 2026); the fees are charged only to takers, the traders whose orders execute against resting limit orders, and fund rebates to the liquidity providers on the other side. Finally, order flow is strikingly one-sided: buys account for 86% of volume in the average market, consistent with participants accumulating positions and holding them to resolution rather than trading in and out of their bets.

Panel B turns to the demand side: most wallets are small and transient, while a thin layer of accounts is extremely active. The median wallet trades in just 4 markets and stakes less than \$10 per trade. Across all of Polymarket’s markets, the median wallet trades 10 markets and \$2,343 of volume (Akey et al., 2026); the wallets in the financial segment are smaller still.

Participation is brief but intense. The median account is active for only 8.3 days between its first and last trade, and a quarter of all wallets complete their entire trading history within about an hour of their first transaction. Conditional on staying, however, traders trade constantly: the median wallet transacts on nearly two thirds of the calendar days spanning its tenure, and the top quarter trades on

every single one of them. Like the market-level flow, wallet-level flow is dominated by purchases: the median account devotes 99% of its volume to buys.

The performance statistics at the bottom of Panel B describe the outcome of this activity for the typical participant. The median wallet wins slightly more than half of its settled positions, yet its cumulative realized profit is negative ($-\$0.05$) and its portfolio return is essentially zero: frequent wins do not imply positive profits. The win-rate distribution reflects the prevalence of very small accounts: a quarter of wallets settle so few positions that they win all of them. The reason lies in the asymmetry of the profit distribution. Gains are typically small and losses large: the 75th percentile of realized profit is $\$0.76$, against a 25th percentile of $-\$14.74$. At the extremes, the bottom percentile of wallets loses more than $\$3,000$, and one wallet in a hundred loses everything it stakes. Mean profit and mean return, moreover, have opposite signs: the mean wallet earns $\$120$ even as the mean portfolio return is -10% , a combination that arises when aggregate trading profits accrue to a small number of large accounts while the typical small account loses a meaningful fraction of what it invests. The same skew characterizes the platform as a whole: across all wallets, [Akey et al. \(2026\)](#) report a median profit of $-\$2$ with the top percentile of winners collecting 76.5% of gains, and [Reichenbach and Walther \(2025\)](#) find fewer than 30% of traders in profit.

Taken together, the two tables describe a platform that expands primarily not by adding new sources of uncertainty but by slicing existing ones ever more finely—shorter horizons and ever-faster reissue cycles—and a population of small, short-lived, buy-and-hold accounts that absorbs each new vintage of contracts. The remainder of this section examines how trading activity migrated across this menu and how the menu has shaped trader performance.

5.3 The Migration of Trading Activity

Section 5.1 documented how Polymarket’s contract offering evolves over time. We now turn to how trading activity evolves in response to these changes. We focus on cryptocurrency markets because new contract structures and issuance frequencies are introduced sequentially, allowing the market response to each innovation to be observed separately.⁶ By contrast, most innovations in equities and

⁶The 4-hour and hourly frequencies are introduced only days apart in cryptocurrency markets. We therefore treat them as a single innovation event throughout the analysis.

commodities are introduced within the same late-2025 period, so changes in structure and frequency overlap and the trading response cannot be cleanly attributed to a particular design margin.

Figure 2 traces the weekly composition of crypto trading volume across contract designs, allowing us to compare the introduction of new contracts with the trading they subsequently attract. The striking pattern is that trading activity shifts almost immediately to the new design. Specifically, panel (a) shows a recurring pattern of rapid volume reallocation across payoff structures. Each newly introduced design quickly absorbs trading from the incumbent structure, leading volume to reconcentrate around the latest innovation.

Hit-price contracts capture roughly one-third of trading in their first week and more than 90% within a month, largely displacing thematic markets. Multi-strike contracts subsequently absorb a substantial share of hit-price trading, reaching about two-thirds of weekly volume at their peak. Up/Down contracts complete this progression: after a slower initial take-up, they steadily erode the market shares of all earlier structures and account for close to 90% of trading by the end of the sample. Neg-risk contracts are the only exception. They were introduced only 35 days before the launch of Up/Down contracts, leaving little time to complete the rapid volume reallocation observed for earlier innovations. Although their market share continued to increase after the introduction of Up/Down contracts, the subsequent Up/Down markets interrupted the transition before neg-risk contracts could become the dominant payoff structure.

Panel (b) reveals a repeated migration of trading toward progressively shorter issuance horizons. Each newly introduced frequency rapidly absorbs trading from the previously dominant one, so that the center of market activity continually shifts toward shorter-lived contracts.

The introduction of monthly contracts triggers the first major reallocation of trading activity, with monthly markets accounting for over 90% of weekly volume within about a month of launch, largely replacing one-off markets. Weekly contracts then become the dominant issuance frequency, capturing roughly three-quarters of trading volume before trading shifts again toward shorter horizons. Daily contracts continue this migration, but are themselves overtaken by the subsequent introduction of intraday markets (4-hour and hourly, combined), which quickly become the primary venue for trading. The process repeats as fifteen-minute contracts absorb much of the remaining intraday volume, before

five-minute contracts emerge as the dominant frequency. Despite existing for only seven weeks, five-minute markets attract about one-fifth of trading volume in their first week and nearly two-thirds by the end of the sample. Consequently, contracts issued more frequently than once per day account for more than half of crypto trading volume by late October 2025 and nearly 90% by the end of the sample.

Overall, trading activity consistently migrates toward newly introduced contract designs. Earlier designs continue to trade throughout the sample, but each successive innovation quickly captures a substantial share of market activity, leaving trading increasingly concentrated in the newest contracts.

Figure OA.2 in the Online Appendix makes the rotation explicit at the trader level. For every launch, we follow the incumbents of each existing design—the traders active in it during the month before—and measure, over the month after, both their take-up of the new design and the change in their presence in the old one. Adoption is strongest exactly where the menu is moving: within a month of the five-minute launch, 58% of fifteen-minute incumbents had already traded the new contract. Exit from the contracts traders were already trading is just as systematic: in every launch window, the number of incumbents still active in the existing design falls, with declines ranging from single digits to more than 90%, and the deepest exits follow the launches of the Up/Down structure and of the fastest frequencies. A newly introduced design, that is, does not add to the set of contracts traders use so much as replace it.

Figure 3 reports the allocation of trading volume over the final seven weeks of the sample, from 12 February 2026 onward—the first date on which every design in the platform’s menu was simultaneously available—so that the shares reflect choice among designs rather than their availability. The resulting allocation is highly concentrated. Up/Down contracts account for 91% of crypto trading volume, and the two fastest issuance frequencies account for 80%, with 59% in five-minute and 21% in fifteen-minute contracts; monthly contracts, which dominated trading in late 2024, represent a negligible share.

Panels (b) and (c) characterize trading behavior across the menu and show that neither the size nor the frequency of trades scales proportionately with the contract horizon. Average trade size declines far less than proportionately as horizons shorten: reducing the contract duration from one month to

five minutes—a factor of approximately 8,600—lowers the average stake from \$99 to \$11, a factor of nine. Stakes are therefore nearly invariant to the horizon, and the capital committed per unit of contract duration is roughly three orders of magnitude larger in the fastest contracts. Small stakes in short-dated contracts are familiar from equity options, where half of retail trades are in contracts with less than a week to expiration and over 40% are for \$250 or less (Bryzgalova, Pavlova, and Sikorskaya, 2023), and from Polymarket itself, where the most and least profitable deciles of traders differ in how often they trade but not in how much they stake per bet (Reichenbach and Walther, 2025). What is distinctive here is that the sequential compression of the horizon lets us observe how the frequency of such positions rises as the opportunities to take them multiply. Although approximately 8,600 five-minute contracts are issued over the lifetime of a single monthly contract, the average number of trades per user per month rises only from 5 to 343, a factor of about seventy, implying that the typical user participates in a small fraction of the available contract cycles. Because the increase in trading frequency substantially exceeds the decline in trade size, the two deviations do not offset: average monthly volume per user is approximately \$3,800 in five-minute contracts, compared with approximately \$500 in monthly contracts.

A similar pattern holds across structures: Up/Down contracts combine small stakes (\$11) with high trading intensity (370 trades per user per month), whereas the richer structures combine larger stakes (\$59–\$90) with at most a dozen trades per user per month. Shorter horizons are thus associated with substantially higher trading volume per user—the margin a volume-maximizing platform has an incentive to expand—and, under the complete menu, the modal transaction is a small, short-horizon binary position in the platform’s simplest contract design.

5.4 Trader Performance

The previous sections documented how the platform expanded its menu and how trading migrated to each new design. We now measure what that migration cost the traders, using the same wallet-level ledger to compute realized performance in cryptocurrency markets separately for each contract design. The evidence on retail products elsewhere is largely negative: yield-enhancement products carry embedded fees large enough to make their expected returns negative (Vokata, 2021), and more

complex structured products carry higher hidden markups and lower ex post returns (C  lerier and Vall  e, 2017).

We measure performance at the level of a wallet–design pair. For each wallet and design, we pool every trade the wallet executes in resolved cryptocurrency markets of that design and compute realized profit as the settlement value of terminal positions, plus revenue from sales, minus the cost of purchases; the realized return divides this profit by the cost of purchases. Dollars are pooled before dividing so that positions offsetting across the two outcomes of the same market net out: a wallet holding one Up share bought at \$0.55 and one Down share bought at \$0.44 earns approximately +1% with certainty, whereas the average of its per-position returns can be markedly negative. Transaction fees are embedded in the cash flows and require no separate adjustment. We report returns rather than dollar profits because the capital deployed varies by orders of magnitude across wallets and designs. We exclude wallet–design pairs with net short positions or missing settlement prices, and we winsorize returns at the 0.01st and 99.99th percentiles of each design’s own distribution so that means are not driven by a small number of extreme ratios. Figure 4 reports, for each design, the mean trader-level return with and without wallet fixed effects, the share of wallets whose realized profit is positive, and the implied speed at which trading capital is depleted.

Panels (a) and (b) report the results by contract structure. In the four richer structures, mean trader-level returns range from -4.5% to -9.4% , yet between 60% and 65% of wallets end with a positive realized profit, indicating that aggregate losses are concentrated in a minority of accounts. Up/Down contracts depart from this pattern: the mean return is -15.0% and only 34% of wallets end in profit. Panels (c) and (d) repeat the exercise by issuance frequency. The share of wallets in profit falls at every increase in issuance frequency: 69% in one-off contracts, 60% in monthly, 57% in weekly, 50% in daily, 47% in four-hour, 40% in hourly, 35% in fifteen-minute, and 25% in five-minute contracts. Raw mean returns follow the same ordering, from -1.9% in monthly contracts to -22.5% in five-minute contracts. Combined with the volume shares in Figure 3, these results imply that trading activity is most concentrated precisely where measured performance is worst: the designs that absorb 91% and 80% of volume are those in which traders’ losses are largest.

Trader-level returns are, however, uninformative about the horizon over which losses accumulate,

because they do not incorporate the rate at which positions are opened. Panels (e) and (f) combine the two ingredients. For each design, we combine the mean loss per dollar staked—taken from the same winsorized wallet–design returns—with the mean number of positions a wallet opens per day over its span in the design, and we compute the number of days before \$100 of trading capital, reinvested position by position at that loss rate, is expected to fall below \$1. Both ingredients deteriorate with speed: from one-off to five-minute contracts, the per-position disadvantage roughly quadruples, from -5.9% to -22.5% , while the number of positions per active day rises by a factor of sixteen, from 0.7 to 11.4. The two margins compound: the expected lifetime of this capital falls from 105 days in one-off and 404 days in monthly contracts to 54 and 48 days at weekly and daily horizons, 44 days in four-hour, 15 days in hourly, 5.2 days in fifteen-minute, and 1.6 days in five-minute contracts; across structures, it falls from 115–167 days in thematic and hit-price contracts to 4.8 days in Up/Down contracts. Faster designs thus deplete traders’ capital through both channels at once: worse odds per position, applied far more often.

Viewed through the framework of Section 4, these results connect the platform’s design choices to their distributional consequences. The contracts toward which the menu, and with it trading volume, migrated are those in which speculative positions turn over fastest, and the turnover that maximizes the volume a trader generates also determines the rate at which the trader’s capital is expected to be depleted. Simpler and faster designs are thus more profitable products for a volume-maximizing platform and worse gambles for the individuals who trade them.

5.4.1 Results Controlling for Fixed Effects

The comparisons above do not account for persistent differences in trader ability, which are well documented among active retail traders ([Seru, Shumway, and Stoffman, 2010](#); [Barber et al., 2014](#)). Because traders select into designs, part of the differences could therefore reflect the composition of wallets rather than the designs themselves. If less skilled traders disproportionately choose Up/Down contracts, Up/Down returns would look worse even if the design itself were neutral. To separate the two, we estimate

$$r_{i,d} = \sum_{k \neq \text{base}} \beta_k \mathbb{1}\{d = k\} + \alpha_i + \varepsilon_{i,d}, \quad (17)$$

where $r_{i,d}$ is the realized return of wallet i in design d , $\mathbb{1}\{d = k\}$ are indicator variables for each design other than the reference, α_i is a wallet fixed effect, and $\varepsilon_{i,d}$ is the error term. The fixed effect absorbs everything constant about the trader—skill, sophistication, risk appetite, capital—so each coefficient β_k is identified only from wallets that trade more than one design and measures how much better or worse the same trader performs in design k than in the reference design. For comparability with the raw means, panels (a) and (c) of Figure 4 report each contrast added to the reference design’s raw mean, so that the raw and within-wallet estimates sit on a common scale: the gap between them is the part of a design’s raw performance attributable to who trades it rather than to the design itself. Standard errors are two-way clustered by wallet and entry day, allowing for correlation both across the designs of the same wallet and across wallets that entered a design on the same day and therefore faced the same market conditions.

For contract structures, where the reference design is thematic contracts, the within-wallet estimates preserve the ordering: they range from -3.3% in hit-price to -11.1% in Up/Down contracts, and every contrast is statistically significant. For issuance frequencies, where the reference design is one-off contracts, the estimates temper the extremes but preserve the pattern: within the same wallet, mean returns decline from approximately zero in monthly contracts ($+0.1\%$) to between -16% and -17% at sub-hourly frequencies, and every contrast is again statistically significant. The deterioration of performance across designs thus persists within the same trader: selection tempers its magnitude but does not generate its ordering.

5.5 Attrition and Replacement

Section 5.4 estimated that in the fastest designs a fixed amount of trading capital is exhausted within days. Their volume can therefore be sustained only by new capital—from losing traders who replenish their accounts or from new entrants who replace them—which is the inflow of the capital-flow condition (Observation 1). A large literature documents this cycle in equity markets: individual investors lose money on average (Barber et al., 2009), the least profitable are the most likely to stop trading (Seru, Shumway, and Stoffman, 2010), and new entrants continually replace them (Barber et al., 2020). This section examines whether the same process accompanies the platform’s shift toward faster contracts.

We first measure how long entrants remain active in each design; we then examine the size of the inflow, the designs through which new accounts enter, and the path by which they reach the fastest contracts.

5.5.1 Attrition

Figure 5 measures retention directly. For each design, we follow the wallets that enter the platform through that design—wallets whose first trade occurs in it—and plot the share still active on the platform as a function of the number of days elapsed since that first trade, over the first thirty days.

Retention declines along both design margins, in the same ordering as returns. Thirty days after entry, 79% of the wallets that entered through one-off contracts are still active, as are 62% of monthly and 52% of weekly entrants; retention then falls to 37% for daily entrants, 24% for intraday entrants, and 12% for five-minute entrants. Across structures, 71% of thematic entrants remain active after thirty days, against 30% of Up/Down entrants. Exits also occur earlier where designs are faster: the one-off and monthly retention curves decline slowly and almost linearly, whereas most exits from the fast designs occur within the first week after entry.

A design that retains 12% of its entrants after a month can hold its trader base constant only through continuous replacement, as the remainder of this section documents: by the end of the sample, roughly half of weekly crypto volume is traded by accounts less than a month old. The attrition curves are consistent with the expected lifetime of trading capital estimated in Section 5.4: for five-minute contracts, an expected lifetime below two days for \$100 implies that few entrants trading at the design’s average intensity can remain active for a month on their initial capital, so the 12% who do must have replenished their accounts, traded far less often than the average entrant, or earned atypical returns. The designs with the fastest capital depletion and the fastest attrition are also those whose remaining users trade most intensely, so aggregate volume can grow while each entering cohort shrinks.

5.5.2 The Inflow of New Accounts

The reallocation in Figure 2 could in principle reflect a market that simply grows with each innovation, so that new designs draw new money while old designs retain their own. Figure 6 shows that this is

not the case. Between March 2024 and March 2026, trade volume rises by a factor of roughly 16,000 while the number of active traders rises by a factor of roughly 3,000, and the two series part ways in mid-2025: from the introduction of intraday contracts onward, volume continues to climb while the count of active traders grows only slowly. The later stages of the menu expansion are therefore not driven by proportional growth in the number of traders: volume per active trader rises about fivefold over the sample.

Slow net growth in the trader count, however, need not imply a stable population of traders: it can conceal substantial replacement if departing traders are continually offset by new entrants. Figure 7 documents this replacement directly, splitting weekly crypto volume between existing clients and new clients, defined as wallets whose first trade on the platform occurred within the previous thirty days. New arrivals always carry a meaningful share of volume, and their share trends upward as the menu accelerates: existing clients account for the large majority of trading through the middle of 2025, but from the launch of fifteen-minute contracts onward the split approaches one half, so that by the end of the sample roughly every other dollar of volume is traded by an account less than a month old. Given the short tenures documented in Table 2—the median wallet is active for just over a week—the population of traders is continually replaced and, on average, becomes younger as the platform’s contracts become faster.

Figure 8 provides the counterfactual counterpart to this evidence. At each contract-introduction date, we freeze the pool of wallets that have traded at least once by that date and track the subsequent daily volume of that closed pool. Each frozen pool is the path the platform would follow if inflow stopped at the freeze date—the $I_t = 0$ counterfactual of the capital-flow condition (Observation 1)—and its volume declines or levels off in every case, while actual volume continues to expand. The gap between each frozen pool and the total, which by construction is generated by wallets entering after the freeze, widens accordingly. By early 2026, total volume approaches \$100 million per day, while even the most recently frozen pool—the wallets active before the five-minute launch—generates less than half that amount. The platform’s expansion therefore cannot be sustained by its incumbent trader base; it increasingly relies on activity supplied by later entrants.

5.5.3 Entry and Migration

The remaining question is where these entrants arrive on the menu and how their activity reaches the fastest designs. A natural conjecture is that the fast designs recruit their own users (Barber et al., 2022; Chapkovski, Khapko, and Zoican, 2026): wallets arrive at the platform to trade five-minute contracts and begin there. Figure 9 shows that this is not the case. Panel A reports the composition of entry under the complete contract menu, measured as the share of first-trade volume accounted for by each design: the sample comprises the wallets whose first-ever trade occurs on or after 12 February 2026—the first date on which every design was simultaneously available, the same window as Figure 3—and each wallet’s first trade contributes its volume to the design in which it occurs. Wallets already active before the window are excluded, so the panel describes genuine newcomers choosing from the full menu.

Entry is concentrated at intermediate issuance frequencies rather than at the fastest. Weekly contracts receive 41% of first-trade volume and hourly contracts 35%, so more than three quarters of entering volume arrives through intermediate frequencies; the two fastest designs together receive only 15% (13% in five-minute and 2% in fifteen-minute contracts), and the slowest frequencies—monthly and one-off contracts—receive almost none, partly reflecting that few monthly and one-off contracts are still being listed by this stage of the sample. The contrast with the allocation of total volume over the same window (Figure 3) is stark: the fifteen- and five-minute contracts that absorb 80% of overall volume attract only 15% of first-trade volume. The same pattern holds along the structure dimension, in attenuated form: Up/Down contracts receive 56% of first-trade volume against their 91% share of overall volume, while multi-strike contracts receive 25% of first-trade volume against a 6% share of overall volume. First trades, in short, occur in slower and richer contracts than the ones in which the platform’s volume sits.

Yet first-trade composition differs sharply from where recent entrants’ subsequent volume ultimately concentrates. Figure OA.1 in the Online Appendix repeats the volume decomposition of Figure 2 separately for new and existing clients. The qualitative pattern is similar in the two groups, but new clients concentrate in the latest designs much more quickly: by the end of the sample, their volume sits almost entirely in Up/Down contracts at the two fastest frequencies, whereas existing

clients retain older structures somewhat longer before converging to the same allocation. Newcomers thus do not begin in the designs where their trading volume ultimately accumulates.

This gap between entry and subsequent allocation points to migration. Panel B measures this migration through the sequence of markets each wallet trades, over the full sample. We define a position as the set of a wallet’s trades in a single market, dated by the first of those trades. For each wallet, we order positions chronologically and record, for each pair of consecutive positions, the design of the earlier and the design of the later one. Because a position collects all of a wallet’s trades in one market, consecutive positions are necessarily in different markets, though not necessarily in different designs. Each row of the heatmap reports the distribution of the next position’s design conditional on the current design, in percent, with each pair of consecutive positions contributing one observation.

Three features stand out. First, design choice is persistent: conditional on a wallet opening a further position, that position is in the same design as the current one in 50% of cases for daily contracts, 92% for five-minute contracts, and 99% for Up/Down contracts. Second, the off-diagonal mass flows almost exclusively toward faster and simpler designs, and it moves to adjacent frequencies rather than jumping across them: one-off traders move next to monthly contracts (39%), monthly traders to weekly (16%), weekly traders onward to daily and the intraday frequencies, hourly traders to fifteen-minute contracts (16%), and fifteen-minute traders to five-minute contracts (7%). Third, the fastest designs are nearly absorbing: transitions back toward slower designs are rare, and rarer still from the fastest frequencies. From the hourly frequency onward, essentially no transitions return to daily or slower frequencies, and a wallet whose current position is in an Up/Down contract places its next position in an Up/Down contract 99% of the time. The one exception is the daily frequency, the only design with substantial transitions in both directions: 24% of its transitions move to weekly contracts and 22% to the intraday frequencies. The transition after which reversals effectively cease is thus not the move from slow frequencies to daily but the move from daily to intraday.

Together, the two panels describe how the fastest designs acquire their traders. New accounts enter at intermediate frequencies, in designs whose horizons still resemble those of conventional investment products; subsequent positions move toward faster contracts; and because the movement is one-directional, the population of the fastest designs is composed of traders who arrived through slower

designs rather than of direct entrants.

This mechanism reconciles several earlier findings. It explains how the fastest designs can dominate aggregate volume (Figure 3) while receiving little entering volume: their users are manufactured internally, and once there, each user trades with the intensity documented in Figure 3—hundreds of positions per month at \$10 stakes. It is consistent with the launch-window evidence of Figure OA.2, in which majorities of incumbents adopt each newly introduced faster design within a month: the transitions in Panel B are the steady-state counterpart of those adoption events. It also sharpens the reading of the attrition estimates above, which classify wallets by the design of their first trade: direct entrants into five-minute contracts—the cohort with 12% thirty-day retention—are a small and selected minority, and most five-minute traders instead arrive by migration, having already remained active through several slower designs.

Part of the measured movement toward faster designs also reflects the evolution of the menu itself: the fastest frequencies did not exist for much of the sample, so a wallet’s progression in calendar time partly tracks the introduction sequence documented in Figure 1. Entry replenishes the trader base, and migration determines where the replenished capital is deployed: familiar, slower products attract the inflow of Observation 1, and subsequent trading moves it toward the designs in which it generates the most volume.

6 Conclusion

We study the contract-design choices of a trading platform and their consequences for the traders who adopt them. In our model, a volume-maximizing platform serving traders with heterogeneous beliefs loads its contracts on the outcome traders disagree about, compresses their maturity, and, through the repetition that short maturities produce, erodes the wealth of its traders even when trading is costless; its volume then depends on a continuing inflow of new capital. The complete record of Polymarket’s financial menu matches these predictions: the menu converged, one design at a time, on five-minute binary bets; trading followed each innovation within weeks; performance and retention deteriorated at every step of the sequence; and the volume that survived was financed by accounts less than a month old, recruited through slower designs and moved toward faster ones.

These findings shift attention from the properties of individual contracts to the platform’s choice of which contracts to make available. The literature on prediction markets and on retail trading has largely studied trader behavior conditional on a given set of securities ([Akey et al., 2026](#); [Della Vedova, 2026b](#); [Barber et al., 2022](#); [Bryzgalova, Pavlova, and Sikorskaya, 2023](#)). Our setting makes the composition of that set an object of analysis: the platform determines the menu, and its incentive to generate trading volume favors designs under which trader capital is staked more often. The menu is therefore a channel through which platform incentives shape trader outcomes.

Our evidence is deliberately confined to Polymarket’s financial contracts, where changes in the menu can be separated from the arrival of outside events and where the settlement value is a public price, so that private information plays a smaller role than in event contracts ([Bartlett and O’Hara, 2026](#)) and trading is closer to the heterogeneous-beliefs setting of our model.

The setting also identifies questions beyond our sample. Whether equities and commodities, which had reached only the daily frequency by March 2026, eventually follow the crypto sequence is observable in public data as listing technology standardizes. As prediction markets enter regulated venues, the design margin we document becomes an explicit object of regulatory review; our results indicate that review conducted contract by contract would not detect the deterioration in trader outcomes, which arises from the menu as a whole. Future research should therefore compare platforms with different regulatory status: Kalshi, in particular, lists its contracts under CFTC regulation ([Bürigi, Deng, and Whelan, 2025](#); [Diercks, Katz, and Wright, 2026](#)), and comparing its menu with the one studied here can shed light on the importance of regulation for the contracts offered to consumers.

References

- Abid, M. N. 2026. The Polymarket Paradox: Manipulation, Whale Concentration, and Predictive Accuracy in the World’s Largest Prediction Market. Working Paper.
- Agarwal, V., P. Ghosh, N. Prabhala, and H. Zhao. 2025. Animal spirits on steroids: Evidence from retail options trading in india. Working paper, September 2025.

- Akey, P., V. Grégoire, N. Harvie, and C. Martineau. 2026. Who Wins and Who Loses In Prediction Markets? Evidence from Polymarket. Working Paper.
- Angelini, G. 2026. The Shape of Macroeconomic Beliefs. ArXiv:2606.30040 [econ.EM].
- Barber, B. M., X. Huang, T. Odean, and C. Schwarz. 2022. Attention-induced trading and returns: Evidence from Robinhood users. *The Journal of Finance* 77:3141–90.
- Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean. 2009. Just how much do individual investors lose by trading? *Review of Financial Studies* 22:609–32.
- . 2020. Learning, fast or slow. *The Review of Asset Pricing Studies* 10:61–93.
- . 2014. The cross-section of speculator skill: Evidence from day trading. *Journal of Financial Markets* 18:1–24.
- Barberis, N., and M. Huang. 2008. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review* 98:2066–100.
- Bartlett, R., and M. O’Hara. 2026. Adverse Selection in Prediction Markets: Evidence from Kalshi. Working Paper.
- Basak, S. 2005. Asset pricing with heterogeneous beliefs. *Journal of Banking & Finance* 29:2849–81.
- Beckmeyer, H., N. Branger, and L. Gayda. 2023. Retail traders love 0dte options... but should they? Working paper, December 2023. SSRN 4404704; verify whether a journal version now exists.
- Ben-David, I., F. Franzoni, B. Kim, and R. Moussawi. 2023. Competition for Attention in the ETF Space. *The Review of Financial Studies* 36:987–1042.
- Boyer, B. H., and K. Vorkink. 2014. Stock options as lotteries. *The Journal of Finance* 69:1485–527.
- Bryzgalova, S., A. Pavlova, and T. Sikorskaya. 2023. Retail trading in options and the rise of the big three wholesalers. *The Journal of Finance* 78. Verify page numbers.
- Budish, E., R. S. Lee, and J. J. Shim. 2024. A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market? *Journal of Political Economy* 132:1209–46.

- Bürge, C., W. Deng, and K. Whelan. 2025. Makers and Takers: The Economics of the Kalshi Prediction Market. Working Paper.
- Chague, F., R. De-Losso, and B. Giovannetti. 2019. Day trading for a living? Working paper. SSRN 3423101.
- Chapkovski, P., M. Khapko, and M. Zoican. 2026. Trading gamification and investor behavior. *Management Science* 72:32–56.
- Cuny, C. J. 1993. The role of liquidity in futures market innovations. *The Review of Financial Studies* 6:57–78.
- Célérier, C., and B. Vallée. 2017. Catering to Investors Through Security Design: Headline Rate and Complexity. *The Quarterly Journal of Economics* 132:1469–508.
- Dai, D., R. Jia, and S. Yu. 2026. Settlement Manipulation in Prediction Markets. ArXiv:2606.31675 [q-fin.TR].
- Della Vedova, J. 2026a. Detecting Informed Trading in Prediction Markets: One Event at a Time. Working Paper.
- . 2026b. Who Profits from Prediction Markets? Execution, not Information. Working Paper.
- Diercks, A. M., J. D. Katz, and J. H. Wright. 2026. Kalshi and the Rise of Macro Markets. Working Paper.
- Dubach, P. D. 2026. The Anatomy of a Decentralized Prediction Market: Microstructure Evidence from the Polymarket Order Book. ArXiv:2604.24366 [q-fin.TR].
- Duffie, D., and M. O. Jackson. 1989. Optimal innovation of futures contracts. *The Review of Financial Studies* 2:275–96.
- Eichengreen, B., G. Viswanath-Natraj, J. Wang, and Z. Wang. 2025. Under Pressure? Central Bank Independence Meets Blockchain Prediction Markets. Working Paper.

- Fabi, M., L. Schönleber, V. Ruffo, and R. Marfè. 2026. Cross-Market Pricing in Prediction Markets: A Comparison with Derivatives. Working Paper.
- Galanis, S., C. A. Ioannou, and S. Kotronis. 2024. Information Aggregation Under Ambiguity: Theory and Experimental Evidence. *Review of Economic Studies* 91:3423–67.
- Gennaioli, N., A. Shleifer, and R. W. Vishny. 2012. Neglected risks, financial innovation, and financial fragility. *Journal of Financial Economics* 104:452–68.
- Glaze, J., J. López-Vilaró, and A. Stephan. 2026. Market Expectations of Narrative Disclosure: Initial Evidence from Earnings Mention Prediction Markets. Working Paper.
- Gomez Cram, R., Y. Guo, T. I. Jensen, and H. Kung. 2025. Financial Prediction Markets: A New Measure of Earnings Expectations. Working Paper.
- . 2026. Prediction Market Accuracy: Crowd Wisdom or Informed Minority? Working Paper.
- Goodman, A., and I. Puri. 2025. Overvaluing simple bets: Evidence from the options market. *Journal of Financial Economics* 172:104140–.
- Heimer, R. Z., and A. Simsek. 2019. Should retail investors’ leverage be limited? *Journal of Financial Economics* 132:1–21.
- Hong, H., and J. C. Stein. 2007. Disagreement and the Stock Market. *Journal of Economic Perspectives* 21:109–28.
- Johnston, E. T., and J. J. McConnell. 1989. Requiem for a market: An analysis of the rise and fall of a financial futures contract. *The Review of Financial Studies* 2:1–23.
- Kumar, A. 2009. Who gambles in the stock market? *The Journal of Finance* 64:1889–933.
- Liu, S. 2026. Wisdom of the crowd or wisdom of the insider? insider trading on prediction markets. Working paper, 29 April 2026. SSRN 6678718.
- Lockhart, K. 2025. Betting on everything. *Boston College Law Review* 66:2175–.
- Manski, C. F. 2006. Interpreting the predictions of prediction markets. *Economics Letters* 91:425–9.

- Mitts, J., and M. Ofir. 2026. From Iran to Taylor Swift: Informed Trading in Prediction Markets. Working Paper.
- Nechepurenko, M. 2026. Empirical Evaluation of Deadline-Resolved Information Leakage on Documented Polymarket Insider Cases. ArXiv:2605.02286 [q-fin.TR].
- Ng, H., L. Peng, Y. Tao, and D. Zhou. 2026. Price Discovery and Trading in Modern Prediction Markets. Working Paper.
- Odean, T. 1999. Do Investors Trade Too Much? *American Economic Review* 89:1279–98.
- Oliphant, P. 2026. The return to imitation: Evidence from copy trading on polymarket. Working paper, April 2026. SSRN 6670318.
- Rabetti, D., J. Shao, and C. Zhang. 2026. Beating the Earnings Game: Why Do Prediction Markets Outperform Professional Analysts? Working Paper.
- Reichenbach, F., and M. Walther. 2025. Exploring Decentralized Prediction Markets: Accuracy, Skill, and Bias on Polymarket. Working Paper.
- Seru, A., T. Shumway, and N. Stoffman. 2010. Learning by trading. *Review of Financial Studies* 23:705–39.
- Silber, W. L. 1981. Innovation, competition, and new contract design in futures markets. *Journal of Futures Markets* 1:123–55.
- Simsek, A. 2013a. Financial Innovation and Portfolio Risks. *American Economic Review* 103:398–401.
- . 2013b. Speculation and Risk Sharing with New Financial Assets. *The Quarterly Journal of Economics* 128:1365–96.
- Sirolly, A., H. Ma, Y. Kanoria, and R. Sethi. 2025. Network-Based Detection of Wash Trading. Working Paper.
- Snowberg, E., and J. Wolfers. 2010. Explaining the Favorite–Long Shot Bias: Is it Risk-Love or Misperceptions? *Journal of Political Economy* 118:723–46.

- Snowberg, E., J. Wolfers, and E. Zitzewitz. 2007. Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections. *The Quarterly Journal of Economics* 122:807–29.
- Tashjian, E., and M. Weissman. 1995. Advantages to competing with yourself: Why an exchange might design futures contracts with correlated payoffs. *Journal of Financial Intermediation* 4:133–57.
- Tsang, K. P., and Z. Yang. 2026. The Anatomy of a Blockchain Prediction Market: Polymarket in the 2024 U.S. Presidential Election. ArXiv:2603.03136 [econ.GN].
- Vokata, P. 2021. Engineering lemons. *Journal of Financial Economics* 142:737–55.
- Wolfers, J., and E. Zitzewitz. 2004. Prediction Markets. *Journal of Economic Perspectives* 18:107–26.




Figure 1: The expansion of Polymarket's contract menu

This figure illustrates how the platform expands its contract menu along two dimensions. Panel (a) shows when each contract structure was first introduced within each asset group. Markers give first appearances, while grey lines span the gap between the earliest and latest asset group to adopt a structure. Panel (b) shows when each issuance frequency was first introduced within each asset group. Dots mark first appearances, while grey lines indicate the gap between the earliest and latest asset group to adopt the same frequency. Panel (c) displays the structure-frequency combinations available in each asset group. Marker size is proportional to the number of markets, and horizontal lines indicate the frequency range covered by each structure. Panel (d) plots every market by listing date and issuance frequency; the black line shows, for each month, the median issuance frequency of that month's newly listed markets, with frequencies ordered from one-off to five-minute. Frequencies are taken directly from the platform's series labels (5m, 15m, hourly, 4h, daily, weekly, and monthly); series without a recurring label are classified as one-off.




Figure 2: Contract structure and issuance frequency shares of crypto trading volume over time

The figure traces the weekly composition of crypto trading volume as the platform's contract menu expands. Panel (a) reports the share of weekly crypto trade volume accounted for by each contract structure; panel (b) reports the corresponding share by issuance frequency. Shares sum to 100% within each week. The sample covers series-covered cryptocurrency markets, classified from the series slug using the taxonomy in Figure 1; trade volume is total traded USDC aggregated to the week in which each trade occurred. Dashed vertical lines mark the launch date of each newly introduced design, labelled with the design. Whereas Figure 3 fixes a single post-go-live window in which every frequency is simultaneously available, this figure shows the full time path along which each new design takes up volume.




Figure 3: Crypto trading under the complete contract menu

Each row of the figure compares contract structures (left column) with issuance frequencies (right column). Panel (a) reports the share of total trade volume, panel (b) the average trade size in USDC, and panel (c) the average number of trades per user per month. The sample includes all cryptocurrency markets listed after 12 February 2026 (100,508 markets; \$4.18 billion in trade volume). The window opens on 12 February 2026 because that is the date five-minute markets began trading, and therefore the first date on which every issuance frequency in the platform's menu was simultaneously available. Restricting the sample in this way ensures that the shares describe the allocation of volume across a complete contract menu rather than one that is still being assembled. Shares in panel (a) sum to 100% within each column. Trade volume is measured as total traded USDC aggregated to the market level. Contract structures and issuance frequencies are classified from the series slug using the taxonomy in Figure 1.




Figure 4: Trader performance across contract designs

This figure reports trader-level performance separately for each contract design. The unit of observation is a wallet–design pair. For each wallet and design, we pool all trades the wallet executes in resolved cryptocurrency markets of that design and compute realized profit as the settlement value of terminal positions, plus revenue from sales, minus the cost of purchases; the realized return divides this profit by the cost of purchases. Both outcomes of a market count as one position, and dollars are pooled before dividing, so that offsetting positions in the two outcomes of the same market net out rather than being recorded as a gain and a loss; transaction fees are embedded in the cash flows and require no separate adjustment. Wallet–design pairs with net short legs or missing settlement prices are excluded, and returns are winsorized at the 0.01st and 99.99th percentiles of each design’s own return distribution. Panels (a) and (c) report two estimates of the mean trader-level returns. The grey estimate is the raw mean return across all wallets that traded the design. The blue estimates include trader fixed effects. In particular, we estimate the regression:

$$r_{i,d} = \sum_{k \neq \text{base}} \beta_k \mathbb{1}\{d = k\} + \alpha_i + \varepsilon_{i,d},$$

where $r_{i,d}$ is the realized return of wallet i in design d , $\mathbb{1}\{d = k\}$ are indicator variables for each design other than the base (thematic in panel (a), one-off in panel (c)), α_i is a wallet fixed effect, and $\varepsilon_{i,d}$ is the error term. Throughout, standard errors are the same wallet clustered by wallet and entry day, where a pair’s entry day is the date of the wallet’s first trade in the design: the same wallet appears in several designs, and wallets entering on the same day share the market conditions of that day. Stars denote significance at the 10%, 5%, and 1% levels. Panels (b) and (d) report the share of wallets in each design whose cumulative realized profit is positive; the dashed line marks 50%, the point at which as many wallets gain as lose. Panels (e) and (f) translate the same returns into the expected lifetime of a fixed amount of trading capital. For each wallet–design pair, we compute the loss per dollar staked, defined as minus the winsorized return, together with the number of positions the wallet opens per day over the span between its first and last trade in the design; both quantities are averaged across wallets and annotated beside each bar. \$100 of trading capital reinvested position by position at loss rate L falls below \$1 after $\ln(0.01)/\ln(1 - L)$ positions; dividing by positions per day converts this figure into calendar days, which each bar reports on a logarithmic axis. Designs are classified from series slugs using the taxonomy of Figure 1.




Figure 5: Attrition from the platform, by design of entry

This figure plots the share of entrants still active on the platform as a function of the number of days since their first trade, separately by the design of that first trade. For each design, the sample comprises the wallets whose first trade occurs in that design; a wallet is counted as active at horizon t if it trades again t or more days after entry. Panel (a) groups entrants by the contract structure of their first trade and panel (b) by its issuance frequency; the four-hour and hourly frequencies are combined. The labels report the share of entrants still active thirty days after entry. Designs are classified from series slugs using the taxonomy of Figure 1.




Figure 6: Trading activity and the timing of contract introductions

The figure places each contract introduction on the path of aggregate trading activity. Two series are plotted: the number of distinct wallets trading in the month, and total trade volume. Because the two are measured in different units, both are expressed as an index equal to 100 in the first month of the sample, March 2024, and drawn on a single axis; the vertical scale is logarithmic. This indexation, rather than a second vertical axis, is deliberate. Two independent scales can be positioned so that any pair of series appears to move together, whereas a common base makes the comparison of growth rates a property of the data. Each launch is marked by a point on both series at its own date, interpolated between the adjacent monthly observations, with a dashed line joining them. Panel (a) marks the launch of each contract structure and panel (b) the launch of each issuance frequency; the two series are identical across panels. The sample covers cryptocurrency series markets, the same universe as Figure 2. The first and last calendar months of the ledger are incomplete and are excluded, leaving twenty-five months from March 2024 to March 2026. Over this period trade volume rises by a factor of roughly 16,000 while the number of active traders rises by a factor of roughly 3,000, so volume per active trader grows about fivefold. The two series move together until the middle of 2025 and diverge thereafter, indicating that the later expansion of the contract menu is associated with more intensive trading by a slowly growing set of participants rather than with a proportional inflow of new ones.




Figure 7: New versus existing clients' share of crypto trading volume

The figure reports the weekly share of crypto trade volume traded by new versus existing clients. A client is *new* in a given week if its first-ever trade in the sample occurred within the previous 30 days, and *existing* otherwise; tenure is measured over all asset groups in the sample. The sample covers series-covered cryptocurrency markets, the same universe as Figure 2. Shares sum to 100% within each week, and dashed vertical lines mark the launch date of each newly introduced design.




Figure 8: Frozen-pool volume: the counterfactual without replacement

This figure tracks the daily crypto trading volume of successively frozen trader pools against actual total volume. Each colored line freezes the trader pool at the introduction date of a new contract design (see Figure 1): all wallets that had traded at least once by that date are retained, and no wallet entering after the introduction contributes to that line. The dark line reports the actual daily volume of all wallets, including those that arrived after any freeze. Both series are smoothed with a 14-day trailing moving average. The gap between the dark line and a frozen-pool line measures the volume attributable to wallets that entered after the corresponding freeze date. The sample covers series-covered cryptocurrency markets, the same universe as Figure 2.




Figure 9: Entry and transitions across contract designs

This figure describes how traders arrive at the platform’s designs and how they move between them. In Panel A, the sample comprises wallets whose first-ever trade occurs on or after 12 February 2026—the first date on which every design in the menu was simultaneously available, the window of Figure 3; each such wallet contributes its first trade, and the panel reports the share of this first-trade volume accounted for by each contract structure (left) and issuance frequency (right). Wallets already active before the window are excluded. In Panel B, computed over the full sample, a position is a wallet’s engagement with one market, dated by its first trade in it; positions are ordered chronologically within wallet, and the heatmaps report, in percent, the design of the next position conditional on the design of the current one, for contract structures (left) and issuance frequencies (right), with each transition counted once and rows summing to 100. Darker cells indicate more frequent transitions; the four-hour and hourly frequencies are combined in Panel B, and designs with fewer than ten distinct markets are excluded from the matrices. Designs are classified from series slugs using the taxonomy of Figure 1.

Table 1: Series Structure by Asset Group

	N Series	N Markets	Avg. Traders	Volume (\$m)	N Tickers	Mkts/Tkr	Avg. Span
<i>Panel A: Crypto (series markets: 210,469)</i>							
<i>A.1 By structure</i>							
Up / Down	52	168,774	407.50	8,544.43	12	14,064.50	1.23
Multi-strikes	17	21,736	118.63	1,571.47	8	2,717.00	3.07
Neg-risk (multi-outcome)	8	15,306	86.00	753.84	4	3,826.50	4.41
Hit price (price reach)	20	4,174	554.28	1,531.70	12	347.83	11.30
Other / thematic	28	479	661.33	116.01	15	31.93	40.35
<i>Total</i>	124	210,469	357.78	12,517.45	27	7,795.15	1.94
<i>A.2 By frequency</i>							
15 min	7	72,113	411.77	3,829.64	7	10,301.86	0.83
5 min	7	61,920	520.57	2,466.22	7	8,845.71	0.98
Hourly	10	34,047	178.04	1,686.06	8	4,255.88	2.31
Weekly	21	19,740	206.99	2,431.67	9	2,193.33	6.98
4 hour	15	18,292	25.95	188.47	7	2,613.14	0.60
Daily	18	2,900	188.50	535.20	9	322.22	1.96
Monthly	22	1,212	1,593.41	1,295.65	14	86.57	29.47
One-time / event	24	245	1,113.21	84.54	15	16.33	73.98
<i>Total</i>	124	210,469	357.78	12,517.45	27	7,795.15	1.94
<i>Panel B: Stocks/Earnings/Indices (series markets: 9,881)</i>							
<i>B.1 By structure</i>							
Multi-strikes	36	4,472	54.42	36.12	23	194.43	7.66
Up / Down	173	2,456	98.40	72.68	164	14.98	2.74
Neg-risk (multi-outcome)	12	1,991	45.89	11.61	12	165.92	7.48
Hit price (price reach)	38	695	98.07	12.81	26	26.73	23.66
Other / thematic	9	267	788.28	71.84	9	29.67	33.56
<i>Total</i>	267	9,881	86.53	205.07	191	51.73	8.22
<i>B.2 By frequency</i>							
Weekly	51	6,151	46.66	41.12	32	192.22	5.97
Daily	26	2,191	84.89	65.04	25	87.64	1.86
Monthly	36	1,096	109.64	19.86	22	49.82	26.83
One-time / event	154	443	591.24	79.05	154	2.88	24.92
<i>Total</i>	267	9,881	86.53	205.07	191	51.73	8.22
<i>Panel C: Commodities/Forex (series markets: 633)</i>							
<i>C.1 By structure</i>							
Up / Down	12	416	32.29	3.08	11	37.82	1.72
Hit price (price reach)	7	127	410.20	51.64	6	21.17	13.64
Other / thematic	7	48	163.88	2.46	5	9.60	38.30
Multi-strikes	6	42	171.38	2.72	4	10.50	35.43
<i>Total</i>	29	633	127.31	59.90	17	37.24	9.12
<i>C.2 By frequency</i>							
Daily	12	416	32.29	3.08	11	37.82	1.72
One-time / event	14	208	322.54	56.81	8	26.00	24.19
Weekly	3	9	7.89	0.01	3	3.00	2.87
<i>Total</i>	29	633	127.31	59.90	17	37.24	9.12

This table reports the structure of the market universe, restricted to markets covered by a series. Each market is classified by contract structure and by issuance frequency based on the naming convention of its series. A market is one binary contract, and volume is realized trade volume in USDC, expressed in millions. Panel A refers to Crypto markets, Panel B to Stocks/Earnings/Indices markets, and Panel C to Commodities/Forex markets. Within each panel, sub-panel 1 classifies markets by structure and sub-panel 2 by frequency, so the two sub-panels partition the same set of markets and share a common *Total* row. For each category, we report *N Series*, the number of distinct series; *N Markets*, the number of markets; *Avg. Traders*, the mean number of unique traders per market; *Volume*, total trading volume in millions of USDC; *N Tickers*, the number of distinct underlying tickers; *Mkts/Tkr*, the ratio of markets to tickers; and *Avg. Span*, the mean created-to-closed lifespan in days.

Table 2: The Platform in Cross-Section: Market and User Characteristics

<i>Panel A: Market characteristics (unit: market; N = 220,983)</i>								
	N	Mean	SD	P1	P25	Median	P75	P99
Trading volume (USDC)	220,983	57,843.44	236,604.17	0.00	3,610.17	13,056.85	49,127.96	565,185.85
Number of trades	220,983	2,771.85	4,801.77	0.00	115.00	946.00	2,595.00	21,277.36
Number of unique traders	220,983	344.99	547.37	0.00	40.00	174.00	353.00	2,173.18
Average trade size (USDC)	213,238	44.01	140.66	2.55	8.11	12.86	28.82	500.00
Buy share of volume	213,238	0.86	0.11	0.57	0.79	0.87	0.94	1.00
Total fees (USDC)	220,983	71.68	199.99	0.00	0.00	0.00	30.66	956.44
Fee rate (fees/volume)	213,238	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Market lifespan (days)	220,983	2.24	6.40	0.13	1.00	1.01	2.08	24.48
<i>Panel B: User characteristics (unit: wallet; N = 1,253,061)</i>								
	N	Mean	SD	P1	P25	Median	P75	P99
Markets traded	1,253,061	60.84	543.90	1.00	1.00	4.00	15.00	1,136.00
Average trade size (USDC)	1,253,061	111.03	650.09	0.50	2.85	9.24	38.56	1,719.10
Buy share of volume	1,253,061	0.81	0.22	0.42	0.56	0.99	1.00	1.00
Total fees paid (USDC)	1,253,061	12.64	728.07	0.00	0.00	0.00	0.00	127.19
Trading tenure (days)	1,253,061	63.45	110.31	0.00	0.04	8.30	70.03	428.70
Active days (% of trading days)	1,253,061	56.81	42.19	0.98	9.23	62.50	100.00	100.00
Win rate	1,249,907	0.56	0.37	0.00	0.28	0.53	1.00	1.00
Realized profit (USDC)	1,249,907	120.28	34,426.10	-3,072.36	-14.74	-0.05	0.76	2,597.97
Portfolio return	1,249,907	-0.10	0.66	-1.00	-0.21	-0.00	0.01	1.02

This table reports cross-sectional summary statistics for the cleaned sample, which comprises Crypto markets together with the Stocks, Earnings, Indices, Commodities, and Forex subcategories of the Economy & Finance category. The sample is restricted to resolved markets that are covered by a series, the same universe reported in Table 1; Panel B accordingly describes wallets aggregated over their trades in those markets. We first compute the value of each variable at the market or wallet level, and then report the distribution of the variable across all markets or wallets. For each variable, we report the number of observations used in the second step of the computations, the mean, standard deviation, as well as the 1st, 25th, 50th, 75th, and 99th percentiles. Panel A refers to market characteristics: *Trading volume*, the total traded volume in USDC; *Number of trades*, the count of executed trades; *Number of unique traders*, the count of distinct wallets active in the market; *Average trade size*, the mean USDC per trade; *Buy share of volume*, the fraction of volume executed as buys; *Total fees*, total fees paid in USDC; *Fee rate*, fees divided by volume; and *Market lifespan*, the number of days from creation to resolution. Panel B refers to user characteristics: *Markets traded*, the number of distinct markets traded; *Average trade size*, the mean USDC per trade; *Buy share of volume*, the fraction of volume executed as buys; *Total fees paid*, total fees in USDC; *Trading tenure*, the number of days from the first to the last trade; *Active days (% of trading days)*, the number of distinct calendar days on which the wallet trades, as a percentage of the whole calendar days spanned by its tenure; *Win rate*, the fraction of settled positions closed at a profit; *Realized profit*, cumulative realized profit in USDC; and *Portfolio return*, the wallet's realized return. Detailed variable definitions are provided in Appendix Table A.1.

A Variable Definitions

Table A.1: Definition of Summary-Statistics Variables

Variable	Definition
<i>Panel A: Market characteristics (unit of observation: market)</i>	
Trading volume	Total cash value of all executed trades in the market.
Number of trades	Number of executed trades in the market.
Number of unique traders	Number of distinct wallets that trade in the market.
Average trade size	Trading volume divided by the number of trades.
Buy share of volume	Share of trading volume executed as buy orders.
Total fees	Sum of per-trade transaction fees, where a trade's fee is the gap between the cash exchanged and the fee-free notional price $\times$ quantity.
Fee rate	Total fees divided by trading volume.
Market lifespan	Number of days from the market's creation to its resolution.
<i>Panel B: User characteristics (unit of observation: wallet)</i>	
Markets traded	Number of distinct markets traded by the wallet.
Average trade size	Trading volume divided by the number of trades.
Buy share of volume	Share of the wallet's trading volume executed as buy orders.
Total fees paid	Sum of transaction fees paid across the wallet's trades.
Trading tenure	Number of days between the wallet's first and last trade.
Active days (% of trading days)	Percentage of trading-tenure days on which the wallet executes at least one trade.
Win rate	Fraction of settled positions with positive realized profits.
Realized profit	Cumulative realized profit across all settled positions. For each position, realized profit equals the settlement value of the remaining holdings minus the net cash invested ($\text{profit} = q_{\text{net}} p_{\text{settle}} - c_{\text{net}}$). Transaction fees are already embedded in cash flows and therefore require no separate adjustment.
Portfolio return	Wallet's realized profit divided by its gross purchase cost.

All monetary values are in USDC. Panel A variables are computed over the full market universe of the sample, so that markets with no trades contribute zero volume and zero trades; Panel B variables are computed over the wallets that trade in the sample.

B Series Coverage




Figure A.1: Series-based versus one-off markets.

For each asset group, the figure reports the share of markets that are issued as part of a series (“has series”) versus one-off markets that carry no series (“one-off”), measured both by the number of markets (left) and by trading volume (right). Shares are computed over the full market universe of the cleaned sample. Series-based markets account for the overwhelming majority of both markets and volume in each group, which motivates our focus on the series structure.

C Model Derivations

C.1 Equilibrium and the Speculative Gain

Substituting the background exposure (2) and the contract (3) into net worth (4) and collecting the coefficients on v_1 and v_2 ,

$$n_g = (e_g - x_g p) + (s_g + x_g c_1) v_1 + (s_g \alpha + x_g c_2) v_2. \quad (18)$$

Writing $\mu_g \equiv \mathbb{E}_g[v_1]$ and using the beliefs and the unit, independent variances in (1),

$$\mathbb{E}_g[n_g] = e_g - x_g p + (s_g + x_g c_1) \mu_g, \quad (19)$$

$$\text{Var}_g[n_g] = (s_g + x_g c_1)^2 + (s_g \alpha + x_g c_2)^2. \quad (20)$$

Substituting into the objective (5), differentiating in x_g , and using $c_1^2 + c_2^2 = 1$ yields the demand

$$x_g = \frac{c_1 \mu_g - p}{\theta} - s_g (c_1 + \alpha c_2); \quad (21)$$

the second-order condition holds because the second derivative equals $-\theta < 0$. Market clearing in zero net supply requires $x_O + x_P = 0$; summing (21) over the two groups and using $\mu_O + \mu_P = 0$ and $s_O + s_P = 0$ leaves $-2p/\theta = 0$, so $p = 0$, and setting $p = 0$, $\mu_O = +D$, $s_O = +1$ yields (6) with speculative part $x_S = c_1 D/\theta$. Under the objective distribution $\mathbb{E}[v_1] = 0$, so $\mathbb{E}[a] = 0$ and any position has objective expected payoff zero: gross of trading costs the contract is a fair, purely redistributive bet.

C.2 The Optimal Contract

From (6) with $p = 0$, $|x_O| = |c_1(D/\theta - 1) - \alpha c_2| = |\gamma \cdot (c_1, c_2)|$ with $\gamma \equiv (D/\theta - 1, -\alpha)$. By the Cauchy-Schwarz inequality this is maximized on the unit circle at $(c_1, c_2) \propto \gamma$, giving (9). Writing $z \equiv D/\theta - 1$, the loading $c_1^* = z/\sqrt{z^2 + \alpha^2}$ has $dc_1^*/dD = \alpha^2/[\theta(z^2 + \alpha^2)^{3/2}] > 0$, so c_1^* is strictly increasing in D and in $1/\theta$, from $-1/\sqrt{1 + \alpha^2}$ at $D = 0$, through zero at $D/\theta = 1$, to one as $D/\theta \rightarrow \infty$, establishing Proposition 1.

C.3 Horizon Comparative Statics

Write (14) as $\varphi(h^*) \equiv D(h^*) - h^*D'(h^*) = F/m$. Since $D(0) = 0$ and D' is strictly decreasing, $D(h) = \int_0^h D'(u) du > hD'(h)$ for every $h > 0$, so $\varphi > 0$; moreover $\varphi'(h) = -hD''(h) > 0$, so φ is strictly increasing and (14) has a unique solution. Total differentiation gives $\partial h^*/\partial F = 1/(m\varphi'(h^*)) > 0$ and $\partial h^*/\partial m = -F/(m^2\varphi'(h^*)) < 0$; since $m = \tau c_1/\theta$, raising the trading cost τ or the loading c_1 lowers h^* , and $F \rightarrow 0$ forces $\varphi(h^*) \rightarrow 0 = \varphi(0)$ and therefore $h^* \rightarrow 0$, establishing Proposition 2. Under the binary specification (11), $D(h) = k\sqrt{h}$ with $k \equiv \sqrt{2/\pi}\mu/\sigma$, so $\varphi(h) = \frac{k}{2}\sqrt{h}$ and $h^* = 4F^2/(mk)^2$: the optimal horizon falls with the square of the listing cost.

C.4 Wealth Concentration

By calendar date T a trader on a horizon- h design has completed $t = T/h$ cycles, so let $Y_t \equiv \ln(W_t/W_0) = \sum_{k=1}^t \xi_k$ with ξ_k i.i.d., mean $\mathbb{E}[\xi_k] = g < 0$ from (15), and finite variance σ_ξ^2 . By the strong law of large numbers $Y_t/t \rightarrow g < 0$ almost surely, so $\Pr(W_t < W_0) = \Pr(Y_t < 0) \rightarrow 1$, giving the majority-loss claim of Proposition 3; $\mathbb{E}[W_t] = W_0$ is constant because wealth is a product of independent multipliers each with unit mean, which supports the conservation statement in the text. The distribution of Y_t depends on the pair (T, h) only through the cycle count $t = T/h$, so halving the horizon halves the calendar date at which any given wealth distribution is reached, which is the rescaling claim. The remaining statements in the text follow from the same construction. Each ξ_k takes the two values $\ln(1 \pm f)$ with equal probability and is therefore symmetric about its mean g ; a sum of independent symmetric random variables is symmetric about the sum of their centers, so Y_t is symmetric about gt , its median equals gt exactly, and the median of W_t is $W_0 e^{gt} = W_0 e^{gT/h}$, falling by half whenever T increases by $h \ln 2/|g|$. By the central limit theorem Y_t is asymptotically normal with mean gt and variance $\sigma_\xi^2 t$, so W_t is asymptotically lognormal with log-standard-deviation $\sigma_\xi \sqrt{t}$; the Lorenz curve of a lognormal with log-standard-deviation σ is $L(q) = \Phi(\Phi^{-1}(q) - \sigma)$, so the top- q wealth share is $1 - L(1 - q) = 1 - \Phi(\Phi^{-1}(1 - q) - \sigma_\xi \sqrt{t}) \rightarrow 1$ for every $q \in (0, 1)$, which supports the winning-minority statement in the text. A per-cycle proportional trading cost multiplies each gross return by $(1 - \tau)$, shifting the drift to $g + \ln(1 - \tau) \leq g - \tau$.

C.5 Reservoir Dynamics

Iterating (16) forward gives

$$K_t = (1 - d)^t K_0 + \sum_{j=0}^{t-1} (1 - d)^{t-1-j} I_j. \quad (22)$$

With constant inflow $I_j \equiv I$ the sum converges and $K_t \rightarrow K^* = I/d$: the pool settles at a level proportional to inflow and inversely proportional to the outflow rate d , and volume is non-decreasing in any period exactly when $I_t \geq dK_t$. With geometrically decaying inflow $I_j = I_0 \lambda^j$, $\lambda < 1$, the solution is a linear combination of λ^t and $(1 - d)^t$ and decays geometrically at the slower of the two rates, establishing the decay claim of Observation 1.

D Online Appendix

D.1 New Clients Adoption of and Exit from Contract Designs








Figure OA.1: Contract designs traded by new versus existing clients

The figure repeats the volume composition of Figure 2 separately for new and existing clients, defined as in Figure 7. The top row splits weekly crypto volume across contract structures and the bottom row across issuance frequencies; the left column restricts to new clients and the right column to existing clients. Shares sum to 100% within each week and panel, and dashed vertical lines mark design launch dates.

(a) contract structures (adoption)**(b) issuance frequencies (adoption)****(c) contract structures (exit)****(d) issuance frequencies (exit)****Figure OA.2: Adoption and exit of incumbent traders around new contract designs**

Both panels of this figure count *traders*. For a new design launched on date D , incumbents are traders who traded an existing design during the month before launch and remained active on Polymarket during the month after; the first condition identifies the relevant user base immediately before the launch, the second excludes traders who had already left the platform. Panels (a) and (b) report adoption: the share of incumbents who traded the new design during the post-launch month. Panels (c) and (d) report exit as the change in the number of incumbents still trading the existing design across the launch, (after – before)/before. Because every incumbent traded the existing design before the launch, this change can only be negative (red); it equals the negative of the share who stopped trading the design. The left column covers contract structures and the right column issuance frequencies. Only transitions from earlier to later designs are shown; blank cells indicate either that a transition was not yet possible or that the incumbent design had no active traders in the pre-launch window. Because the 4-hour and hourly designs launched within ten days of one another and their windows almost fully overlap, they are combined into a single intraday category.