

Settlement Manipulation in Prediction Markets

David Dai
Stanford University

Ruizhe Jia
Stanford University

Shihao Yu
Singapore Management University

**SIM KEE BOON
INSTITUTE
FOR FINANCIAL
ECONOMICS**

**LEE KONG CHIAN
SCHOOL OF BUSINESS**

Settlement Manipulation in Prediction Markets

David Dai* Ruizhe Jia[†] Shihao Yu[‡]

Current version: June 30, 2026

*Department of Management Science and Engineering, Stanford University; htdai@stanford.edu.

[†]Department of Management Science and Engineering, Stanford University; ruizhe@stanford.edu.

[‡]Lee Kong Chian School of Business, Singapore Management University; shihaoyu@smu.edu.sg.

Settlement Manipulation in Prediction Markets

Abstract

Prediction markets increasingly list contracts settling on an asset price that holders can move by trading the underlying. We build a model showing that such contracts transfer wealth from prediction-market liquidity traders to manipulators and harm price discovery in the underlying, even as it becomes more liquid. After the launch of Polymarket's five-minute Bitcoin contract, settlement-time spot order flow spikes, causing large price reversals after settlement. Manipulators capture a large amount of profit, mostly from retail. Manipulation is largely absent in the fifteen-minute contracts: lengthening the contract horizon removes it, providing the market-design remedy our model and evidence support.

1 Introduction

Prediction markets have grown with remarkable speed: combined monthly volume on Kalshi and Polymarket rose nearly fivefold in seven months, from under \$5 billion in September 2025 to about \$24 billion by April 2026.¹ Surprisingly, the fastest-growing of these contracts are not written on elections or sports, but on asset prices. Each is a binary contract: it pays \$1 if the asset is above a chosen level at the contract's close, and \$0 otherwise. These began in cryptocurrencies, where they are already large. On Kalshi alone, volume passed \$1 billion a month for the first time in March 2026.² Now they are reaching mainstream finance: Nasdaq and Cboe have filed to list them on equity indices.³ And more importantly, the contract horizon keeps shrinking: from a day, to an hour, to fifteen minutes, and on Polymarket to just five minutes.

Regulators have not merely permitted prediction contracts; they describe their benefits to the public. The CFTC tells investors that “a citrus farmer might buy a weather event contract to hedge against losses that might be caused by a sudden freeze,” and that prediction markets “can sometimes forecast event outcomes better than polling or other forms of forecasting.”⁴ Academic work offers concrete support: the Iowa Electronic Markets forecast U.S. presidential vote shares to within about 1.5 percentage points, outperforming large-scale polls (Wolfers and Zitzewitz, 2004), and a prominent group of economists argues that prediction markets deliver “a lower prediction error than conventional forecasting methods” (Arrow et al., 2008). Even the worry that someone might manipulate the price has largely been set aside: a manipulator can increase price accuracy (Hanson and Oprea, 2009), and documented manipulation attempts reverted quickly (Rhode and Strumpf, 2008). If that logic carried over to asset prices, these contracts would make financial markets more informative, not less.

This paper shows the opposite for asset-price contracts: rather than making the underlying

¹Pew Research Center, “Trading Volume on Prediction Markets Has Soared in Recent Months” (2026).

²CF Benchmarks, “Kalshi Leads Surging Crypto Event Contract Market” (2026).

³Nasdaq and Cboe each filed with the SEC in March 2026 to list binary (fixed-payout) options on equity indices.

⁴Commodity Futures Trading Commission, “Understanding Prediction Markets and Event Contracts” (investor advisory), [cftc.gov](https://www.cftc.gov).




Figure 1. The footprint of settlement manipulation. This figure plots Binance BTC spot activity around the five-minute contract’s close, in the post-launch period. *Left:* absolute order flow per five-minute cycle, folded onto the rolling clock; the dotted vertical line marks the final ten seconds before the close. *Right:* the post-settlement reversal slope γ_j at each lag j around the close, the slope of the lag- j ten-second return on the final-ten-second return. Section 6 details both constructions and establishes the patterns against pre-launch and across-moneyness benchmarks.

market more informative, they degrade its price discovery and transfer wealth from ordinary traders to those who manipulate them. The vulnerability is structural. An asset-price contract settles on a financial price, and that price can be moved by trading the underlying market itself, the cash-settlement manipulation of Kumar and Seppi (1992), now inside a retail prediction market with horizons as short as five minutes. U.S. law already requires every listed contract to be “not readily susceptible to manipulation”;⁵ we show that asset-price contracts raise exactly the concern that standard is meant to prevent. The footprint is in the data (Figure 1): after the five-minute contract went live on Polymarket, in the final seconds before settlement, Binance spot order flow spikes, resulting in price movements that revert shortly after settlement. Clearly, this is a transitory push to manipulate the spot price, not trading on information.

We develop a sequential game-theoretic model of the two markets. The spot market has its usual players (an informed trader who knows the asset’s value and liquidity traders who do not) and now a manipulator, who has bought the contract and trades the spot just before settlement. Her order

⁵Commodity Exchange Act, Designated Contract Market Core Principle 3; reaffirmed for event contracts in CFTC Staff Advisory 26-08 (2026).

pools with the rest, and because she trades on no information, a late buy is now less likely to be the informed trader's. It therefore carries less information about the asset, so the settlement price reveals less of it and price discovery falls. Yet, by lowering the adverse selection the spot maker faces, it leads the spot maker to supply more liquidity, so the manipulation makes the spot market more liquid even as it makes the price less informative.

The same inability to tell who is trading also makes the liquidity traders in the prediction market earn worse-than-fair odds. A contract buy is public, but the spot maker cannot tell the manipulator from an ordinary contract buyer. It therefore puts some probability on a manipulator behind the trade, one who will push the spot toward her side at settlement. Anticipating that push, the spot maker shades its settlement-window quote against it, and the settlement tilts against the side the contract buyer took. A prediction-market liquidity trader who bought with no intent to manipulate is overpaying all the time for this contract. More importantly, we show that a prediction-market maker earns a rent as long as the bid and ask cannot cross. The prediction-market liquidity trader's loss is thus a transfer to the prediction-market maker and the manipulator.

These harms are not inevitable: both the evidence and the model point to one design lever, the contract's horizon. The model explains why the horizon matters. A short horizon leaves little room for the asset's own price discovery before settlement, so a fixed spot push is more likely to determine the contract's resolution. A longer horizon aggregates more ordinary spot-market information before the close, making the price paths in which a push is pivotal rarer and reducing the price-discovery loss on net. The remedy is direct: lengthen the horizon. In the data, the manipulation signature is sharp for the five-minute contract, but much attenuated for the fifteen-minute contract.

Empirically, we study Polymarket's Bitcoin five-minute up/down contract: a binary claim that pays \$1 if Bitcoin is higher at the close of a five-minute window than at its open, and \$0 otherwise. Every trade is recorded on a public blockchain, so each trade and trader is traceable, a transparency centralized exchanges and options venues do not offer. The contract resolves not on Polymarket but against a Chainlink oracle that averages the price across major spot exchanges. The price on Binance, the largest crypto exchange by volume, is an economically tight proxy for that benchmark:

its mid sits about two and a half basis points from the oracle and moves essentially one-for-one with it within each five-minute window, finishing on the same side of the strike as the resolution about 85% of the time, so a push that drives the Binance mid a few basis points past the strike reliably carries the resolution. We therefore measure the settlement-time flow on Binance. The contract's launch on February 12, 2026 offers a clean natural experiment: no five-minute contract existed before that date, so we compare three regimes: no such contract (P1), the 15-minute and 4-hour contracts only (P2), and the five-minute contract live (P3).

As soon as it launched, a relatively small prediction market began redirecting large spot orders to the settlement seconds. In the final ten seconds before each close, trading on Binance spikes: the magnitude of net order flow jumps, and volatility rises with it. The timing is no accident: the spike appears only after the five-minute launch (about 50% above the pre-launch level), and is significantly attenuated at the fifteen-minute horizon. The spike is sharpest where a push is pivotal: in the roughly 6% of cycles whose contract price still implies a near-even outcome just before the close, the near-settlement order-flow jump is about 3.9 times that in the rest. The clearest sign that this is not information is the reversal: within ten seconds the price reverts, by about a quarter in the near-even cycles and a tenth in the others. Real information would persist; the price impact of a manipulative push reverts. Last, the liquidity evidence is asymmetric across venues: near-even closes add depth on Binance, exactly where the model predicts an uninformed push should lower spot adverse selection, while Polymarket depth falls as binary-market makers face direct exposure to a flipping settlement.

We then identify the cycles most likely to contain a manipulative push. For each cycle we measure the size of the final-ten-second net order flow on Binance relative to the cycle's own typical flow. We classify the cycles in the top decile of this measure, about 1,600, as manipulated. Away from the close they are quieter than average, and they fall disproportionately in thin-liquidity hours: 56% in the overnight window and 44% on weekends, against 40% and 27% of normal cycles, where a given dollar of near-settlement flow moves the spot price, and therefore the resolution, the most.

Within these manipulated cycles, the pushes do more than cluster at the close: they change

which side wins. In cycles where the contract was priced near even, a push against the favored side flipped the winner 65% of the time, against 41% in ordinary trading. It works by dragging the Binance price, and thus the oracle price, across the strike. The stronger test is the cycles that were all but decided: even when the market gave one side a 90-to-100% chance just before the close, a push against it reversed the outcome 34% of the time, against 1% in cycles with no push. A bet the market treated as near-certain was overturned one time in three.

We next ask who gains from that and who pays for it. Because Polymarket settles on a public blockchain, every wallet's profit is visible, and we identify as manipulators the traders who gain in the manipulated cycles, an identification based on realized profits rather than a direct observation of intent. Just 821 of them fit, about one in three hundred of the roughly 243,000 who traded the contract, and they take \$8.2 million in the pushed cycles while breaking even in the rest: the edge lives only where the manipulation does. Who pays divides just as sharply. Setting aside the market makers, who quote passively and end each cycle flat, almost none of the loss lands on them: 93% of it falls on retail. These retail accounts are the liquidity traders of the model, and the contract treats them accordingly: they are on the losing side in 65% of manipulated cycles, against 48% of normal ones. All of it comes from a single contract, on a single platform, over two months.

The natural innocent explanation is hedging: a prediction-market maker who has sold the binary contract offsets its risk by trading the spot. Two facts rule that out, and the clearer one is the state of the contract when the pushing happens. A binary contract's sensitivity to the spot price is large only when the price sits near the strike; once one side is almost certain to win, the contract barely moves with the spot, and the prediction-market maker is left with almost nothing to hedge. Yet those nearly decided cycles (where the market already gave one side a 90-to-100% chance) are exactly the ones a push still flips, one time in three. A trade placed when there is nothing to hedge cannot be a hedge; whatever hedging does occur can only come afterward, once the push has dragged the price back toward the strike and revived the exposure. The second fact is the timing. Even when there is something to hedge, a prediction-market maker would build its spot position gradually, as that exposure accumulates over the cycle. Instead the flow arrives in a single burst at the very

end: essentially all of the signed trading falls in the final fifty seconds, and almost none before. A position assembled only in the closing seconds is not a hedge accumulated over the cycle; it is a trade placed to move the settlement price.

The paper proceeds as follows. Section 2 positions the contribution. Section 3 describes the contracts and how they settle. Sections 4 and 5 develop the model and its predictions. Section 6 presents the testable predictions, the data, and the empirical results, and Section 7 discusses the market-design implications.

2 Related Literature

This paper joins three strands of literature that have largely developed apart: the economics of prediction markets, the manipulation of cash-settled derivatives, and the market microstructure of decentralized and crypto venues.

Prediction markets. Prediction markets are valued as information aggregators whose prices forecast events as accurately as expert judgment (Wolfers and Zitzewitz, 2004), a view recently reaffirmed for macroeconomic contracts on Kalshi (Diercks, Katz, and Wright, 2026). A long-standing concern is that such prices are vulnerable to manipulation, yet the literature has largely concluded that manipulation need not impair price accuracy. In a Kyle-style model, a manipulator's order flow raises the returns to informed trading and can even improve price accuracy (Hanson and Oprea, 2009); in laboratory markets, other traders anticipate and discount manipulators' biased orders, leaving accuracy intact (Hanson, Oprea, and Porter, 2006); and documented manipulation attempts in the field have produced only transient, mean-reverting price effects (Rhode and Strumpf, 2008). A growing empirical literature exploits transaction-level data from Polymarket and Kalshi: one strand examines who profits and who is informed, finding that a small group of sophisticated traders captures most of the gains while retail participants fund them (Akey, Grégoire, Harvie, and Martineau, 2026; Gómez-Cram, Guo, Jensen, and Kung, 2026; Della Vedova, 2026); a parallel

strand measures adverse selection and order flow (Bartlett and O’Hara, 2026; Tsang and Yang, 2026; Dubach, 2026) and screens for suspicious accounts (Mitts and Ofir, 2026).

To our knowledge this is the first study of ultra-short, oracle-settled prediction markets on an asset’s price, and the first to identify settlement manipulation in them: a trader who holds the binary contract trades the underlying asset near settlement to move the price on which the contract pays. We show that this manipulation not only transfers wealth from liquidity traders to manipulators in the prediction market, but also harms price discovery in the underlying spot market.

Manipulation of cash-settled derivatives. At its core, our mechanism is the cash-settlement manipulation of Kumar and Seppi (1992): a trader takes a derivative position and then trades the underlying spot to move the cash-settlement price, profiting when the gain on the position exceeds the cost of the spot trade. Another theory paper studying manipulation in derivatives markets is Zhang (2022), in which manipulators can take endogenous derivative contract positions and leave every agent worse off. On the empirical side, the literature documents settlement and expiration distortions in equities and volatility benchmarks (Ni, Pearson, and Poteshman, 2005; Comerton-Forde and Putniņš, 2011; Griffin and Shams, 2018; Baltussen, Terstegge, and Whelan, 2023). Spatt (2014) surveys a variety of manipulation strategies. Finally, a design literature asks how to build benchmarks that resist manipulation (Duffie and Dworczak, 2021).

Our contribution relative to this work is twofold. First, on the theory side, we build a model of cash-settlement manipulation in a new market: ultra-short, oracle-settled prediction markets on an asset’s price. This model is based on the sequential trading framework of Glosten and Milgrom (1985) and on the two-venue price-discovery structure of Zhu (2014). In this market, manipulation harms liquidity traders and degrades the underlying spot’s price discovery. We further provide a market-design remedy: a longer contract horizon lowers both the manipulator’s rents and the price-discovery loss. Second, on the empirical side, because this market settles on-chain—unlike the equity, futures, and volatility-benchmark venues where cash-settlement manipulation is otherwise studied—we observe trading at the level of pseudonymous wallets, which allows a more granular

analysis: we identify the likely manipulators and measure the wealth they extract.

Decentralized and crypto market microstructure. Finally, the paper contributes to the microstructure of decentralized and crypto markets by studying a new on-chain venue: a binary, oracle-settled prediction market. One strand studies the microstructure, price discovery, and quality of crypto trading venues, both centralized and on-chain (Makarov and Schoar, 2020; Lehar, Parlour, and Zoican, 2022; Klein, Kozhan, Viswanath-Natraj, and Wang, 2023; Aoyagi and Ito, 2025; Easley, O’Hara, Yang, and Zhang, 2026; Barbon and Rinaldo, 2026). A second studies liquidity provision and price discovery on decentralized exchanges and automated market makers (Hasbrouck, Rivera, and Saleh, 2024; Capponi and Jia, 2025; Lehar and Parlour, 2025; Capponi, Jia, and Yu, 2026; Hasbrouck, Saleh, and Rivera, 2026). A third studies value extraction through transaction ordering and front-running, and the design flaws and responses it has prompted (Daian et al., 2020; Park, 2023; Capponi, Jia, and Wang, 2025). A fourth concerns on-chain settlement and the price oracles that feed external data to smart contracts (Chiu and Koepl, 2019; Cong, Prasad, and Rabetti, 2025). Our mechanism is nonetheless distinct from these on-chain frictions: Polymarket runs a central limit order book, and the manipulation operates on the centralized spot exchange where the contract settles, working through trading into the settlement window rather than transaction reordering or a compromised oracle.

3 Institutional Background

3.1 Prediction markets and asset-price contracts

A prediction market lets participants trade contracts that pay off on the realized outcome of a future event, so the contract’s price can be read directly as the market’s estimated probability that the event occurs. The established contracts settle on real-world events (elections, policy decisions, sporting results) and are valued as forecasters whose prices have matched or beaten conventional methods (Wolfers and Zitzewitz, 2004; Arrow et al., 2008). One desirable feature stands out: the settlement

event lies outside the traders' control (it is *exogenous* to them).

The fastest-growing asset-price contracts are the cryptocurrency “up/down” markets: a yes-or-no bet on the direction of a reference asset (most heavily Bitcoin) over a fixed horizon, paying off only if the asset ends the horizon above where it began. The horizons are short, from minutes to hours, and the contracts open and settle automatically on a rolling schedule, so a single asset generates a continuous stream of fresh binary markets through the day.⁶

The same structure is now reaching mainstream finance. In March 2026 Nasdaq filed with the U.S. Securities and Exchange Commission to list binary options on the Nasdaq-100 and Nasdaq-100 Micro indices, and Cboe is pursuing comparable “all-or-none” contracts on financial and economic events; like the crypto up/down markets, each pays a fixed amount if a stated price condition is met at expiry and nothing otherwise.⁷ These proposals would carry the asset-price contract onto the most heavily traded U.S. equity indices, so the settlement vulnerability this paper documents in crypto would extend to far larger markets.

3.2 Venues

These contracts trade on venues that differ sharply in design. *Polymarket* is a decentralized, blockchain-based market: its contracts are created and settled automatically by code that runs on a public blockchain (on-chain smart contracts), every trade is recorded on that blockchain, and positions are held in a token pegged to the U.S. dollar (a dollar stablecoin) (Tsang and Yang, 2026; Dubach, 2026). It lists the widest menu of horizons, including five-minute, fifteen-minute, and four-hour up/down contracts on Bitcoin and several other coins. *Kalshi* is a centralized, federally regulated U.S. exchange (trades route through one operator rather than a blockchain) (Diercks, Katz, and Wright, 2026; Commodity Futures Trading Commission, 2026); over our sample it lists only a

⁶Within months of launch, Polymarket's five- and fifteen-minute crypto up/down markets traded over \$4 billion cumulatively and roughly tripled platform daily volume; sources: Polymarket and Chainlink launch announcements, 2026.

⁷Nasdaq's proposed contracts are priced between one cent and one dollar on the Nasdaq-100 and Nasdaq-100 Micro indices (SEC rule-change filing, March 2026); Cboe has announced “all-or-none” event contracts targeting a 2026 launch, subject to regulatory approval. Sources: Nasdaq and Cboe SEC filings and announcements, 2026.

fifteen-minute crypto up/down contract, with no five-minute or four-hour equivalent; its hourly and daily crypto series are price-range markets, not binary up/down.

3.3 Polymarket

Each Polymarket crypto up/down contract pays \$1 if the reference asset's price at the close exceeds its price at the open, and \$0 otherwise. The contracts are created and settled programmatically on a rolling schedule, so a fresh five-minute Bitcoin contract opens every five minutes around the clock. The three horizons trade simultaneously and overlap on the clock (every fifteen-minute contract spans three consecutive five-minute contracts), a feature we exploit below.

Resolution. The five-minute, fifteen-minute, and four-hour contracts settle against a price produced by a decentralized oracle (an outside service that reports a reference price by combining data from many exchanges) rather than against any single exchange. The oracle, Chainlink, aggregates trade data from many centralized and decentralized venues (including the largest spot exchanges) and reduces them to one consensus print through a multi-stage procedure (Cong, Prasad, and Rabetti, 2025; Aspembitova and Bentley, 2022): several professional aggregators each compute a volume-weighted average price (a price that weights each trade by its size) across the venues they cover, independent node operators take the middle value (the median) across aggregators, and the network publishes the middle value across nodes.⁸ The strike against which a contract resolves is the oracle price sampled at the contract's open, and the settlement value is the oracle price at its close; the "up" side wins if and only if the close print exceeds the strike. (The one-hour and daily contracts, which fall outside our sample, instead reference a single exchange's spot price directly.) This aggregated construction is central to the manipulation question we study. Because the resolution price combines many exchanges into one agreed figure, moving it appears to require moving the price on many exchanges at once rather than on just one. Yet within a five-minute

⁸For the fast markets Polymarket reads Chainlink Data Streams, a low-latency pull-based feed, rather than the standard on-chain push feed, so the settlement value can be sampled at the exact close; see Chainlink Data Feeds documentation, <https://docs.chain.link/data-feeds>.

horizon, arbitrage (traders profiting from price gaps across exchanges) keeps the largest single venue and the combined price within a few basis points (hundredths of a percent) of each other. A spot-price move concentrated on one venue can therefore still push the combined price across the strike if the move is large enough relative to that narrow margin. Appendix D measures this gap and the residual disagreement directly.

Product rollout. The three Bitcoin contracts launched in sequence, and that staggered introduction is the backbone of our between-period comparison. The fifteen-minute contract began trading on October 9, 2025 and the four-hour contract on October 15, 2025; the five-minute contract was created on-chain in December 2025 but sat dormant, with no trading, until its public launch on February 12, 2026.⁹ We therefore partition the sample into three periods: a pre-launch baseline before any of these contracts traded, an intermediate period with the fifteen-minute and four-hour contracts live, and a final period with all three (including the five-minute contract) trading together.

4 Theoretical Setup

Two markets operate in parallel. On a spot market a risky asset changes hands; on a separate venue a binary contract settles on that asset's price at the close of trading, with some noise. The settlement rule is what links them: the contract settles on a price the spot market itself discovers—and a trader who holds the contract can move that price by trading the asset. An uninformed manipulator turns this into a strategy. She takes a position in the contract, then trades the spot in the final moments before settlement to push its price toward the side she holds.

4.1 The Two Markets and the Contract

This section describes the two assets the model needs—a risky spot asset and a binary contract written on it—and the rule that settles the contract.

⁹Launch dates are identified from the on-chain trade record; see Section 6.2.

The first asset is the spot asset itself. The fundamental value of the spot asset is a random variable V with equal probability of being $-\sigma$ or $+\sigma$, drawn by nature at stage 0. $\sigma > 0$ is the volatility of the fundamental value and exogenous. One informed spot trader observes the realization. Everyone else holds the symmetric prior. Spot positions ultimately liquidate at the value of V .

The second asset is the binary contract—a claim on the asset’s direction. Crucially, it does not settle on the realized value V but on the spot market’s *posterior mean* of V at the end of the prediction horizon, that is, the value the spot market’s own trading has revealed, perturbed by a small independent noise. This is the link the manipulator exploits: by moving the closing spot price, she moves the very quantity her contract settles on. Write h for the public order flow and $\bar{V}_h$ for the spot market’s posterior mean of V given h (both made precise in Section 4.5).

For a final price path h , the binary contract pays

$$D_\varepsilon(h) = \begin{cases} 0, & \bar{V}_h + \varepsilon\xi < 0, \\ \frac{1}{2}, & \bar{V}_h + \varepsilon\xi = 0, \\ 1, & \bar{V}_h + \varepsilon\xi > 0, \end{cases} \quad (1)$$

where ξ is a random noise term uniformly distributed on $[-1, 1]$, independent of V and all orders, and $\varepsilon > 0$ is an exogenous parameter that controls the noise level.¹⁰

The payoff in (1) makes the contract a *binary* (or digital) option on the spot price: it pays out only when the settlement value $\bar{V}_h + \varepsilon\xi$ clears a fixed threshold, here zero. We therefore adopt standard option terminology throughout. We call that threshold the *strike*; the contract finishes *in-the-money* when the settlement value is above the strike (it pays 1), *out-of-the-money* when below it (it pays 0), and *at-the-money* when the price sits at the strike, where the outcome is still in doubt.

Why settle on a noisy price rather than on the spot price directly? In practice the settlement

¹⁰In the no-noise limit $\varepsilon \downarrow 0$, the expected payoff jumps when $\bar{V}_h$ crosses zero and we might not have an equilibrium. The noise smooths the payoff and guarantees the existence of an equilibrium. In the large-noise limit $\varepsilon \uparrow \infty$, the binary contract decouples from the spot market and its expected payoff converges to $1/2$ regardless of the price path, which makes manipulation unprofitable.

reference is not the manipulator's to dictate. The oracle that resolves the contract reads many exchanges at once—not just the one venue she can trade—and her push lands amid a stream of other traders' orders she does not control. A push therefore carries the settlement across the strike only with some probability, never for sure. The term $\varepsilon\xi$ captures exactly this residual uncertainty, and ε sets how much of it remains.

Let H denote the cumulative distribution function of ξ . The expected payoff of the PM contract given h is

$$\mathbb{E}[D_\varepsilon(h) | h] = H\left(\frac{\bar{V}_h}{\varepsilon}\right) = \begin{cases} 0, & \bar{V}_h \leq -\varepsilon, \\ \frac{1}{2} + \frac{\bar{V}_h}{2\varepsilon}, & -\varepsilon < \bar{V}_h < \varepsilon, \\ 1, & \bar{V}_h \geq \varepsilon. \end{cases} \quad (2)$$

That is, the contract pays 1 if the final posterior mean is weakly above ε , 0 weakly below $-\varepsilon$, and interpolates linearly on the band of width 2ε around zero.

4.2 Timeline

The model has five types of participants—market makers on the spot and prediction-market venues, an informed spot trader, a manipulator, and liquidity traders—who act over five stages (Figure 2). Each order below is publicly observed. At each trading stage the venue's market maker has posted bid and ask quotes, and the arriving order executes against them. Later, Section 4.3 specifies how each participant behaves strategically; here we only lay out the sequence.

Stage 0 (Nature): Draws $V \in \{-\sigma, +\sigma\}$, observed only by the informed spot trader.

Stage 1 (Prediction market): One order Y_1^P arrives on the prediction market (PM), from the manipulator or a liquidity trader.

Stage 2 (Spot market): One order Y_2^S arrives on the spot market, from the informed trader or a liquidity trader.

Stage 3 (Spot market during settlement window): One order Y_3^S arrives on the spot market just

before the close: the manipulator if she entered at stage 1, otherwise an ordinary spot order like stage 2's.

Stage 4 (Settlement): The binary contract settles on the spot market's posterior mean of V plus a small noise after stage 3 ends. All spot positions liquidate at the realized value of V .




Figure 2. The five-stage timeline. Filled dots mark the manipulator's two decision points: entry on the prediction market (Stage 1) and the settlement-window spot trade (Stage 3).

4.3 Players

Market makers. A competitive, risk-neutral market maker stands on each venue and sees the same public information $\{\mathcal{F}_t\}$. Competition drives its expected profit to zero, so it quotes a break-even price—a one-unit bid and a one-unit ask equal to the asset's expected value given the order it faces. One qualification is required: the book cannot be crossed, so the ask never falls below the bid. Where the break-even quotes satisfy this, the maker posts them and earns nothing; where they would cross, it cannot—and the tightest feasible book, a zero spread, lets it earn a rent instead. It cannot distinguish informed or strategic orders from liquidity trades, so the quote conditions on the order alone. The spot maker's quotes are (B_t^S, A_t^S) : at stage 2, (B_2^S, A_2^S) ; at stage 3, after stage-2 outcome h_2 , (B_{h_2}, A_{h_2}) . The prediction-market maker's stage-1 quotes are (B_1^P, A_1^P) .

Informed trader. The informed trader knows the realized V and trades to profit from it. She arrives with probability $\pi \in (0, 1)$ at each *ordinary* spot stage—stage 2 always, and stage 3 whenever the manipulator has not entered—and follows the Glosten–Milgrom rule (Glosten and Milgrom, 1985),

$$\text{buy if } V > A_t^S, \quad \text{sell if } V < B_t^S, \quad \text{stay out otherwise.} \quad (3)$$

That is, she buys when the asset is worth more than the ask and sells when it is worth less. Because break-even quotes lie inside $[-\sigma, \sigma]$, one side always pays, so she never stays out.

Liquidity traders. Liquidity traders are uninformed and trade for reasons outside the model, buying or selling at random. One sits on each venue: on the prediction market a liquidity trader takes the stage-1 slot whenever the manipulator does not enter, and on the spot market a liquidity trader takes any ordinary spot slot the informed trader does not. Either way exactly one order prints per stage, so the stage-2 outcome is a buy or a sell.

Manipulator. The manipulator never observes V . She profits not from information but from the contract itself: by trading the spot in the settlement window she moves the price the contract pays on. At stage 1 she enters the prediction market with probability $p_M \in [0, 1]$, buying or selling the contract; with probability $1 - p_M$ a liquidity trader takes the stage-1 slot instead. The prediction-market maker sees only the order, not who sent it, so it cannot tell an entering manipulator from a liquidity trader. Conditioning on a stage-1 buy (the sell branch is symmetric), let $M_t(\delta) = \mathbb{P}(\Delta_t = \delta \mid \mathcal{F}_t)$ be the public belief that her contract inventory after stage t is $\delta \in \{-1, 0, +1\}$:

$$M_1(-1) = 0, \quad M_1(0) = 1 - p_M, \quad M_1(+1) = p_M.$$

She does not trade at stage 2, so $M_{3-} = M_1$. What happens at stage 3 depends on whether she entered. If she did, she alone takes the settlement-window slot and chooses how hard to push: with probability $\alpha \in [0, 1]$ she buys, moving the close toward the side of the contract she holds, and otherwise does nothing. This manipulation probability is her second choice variable, alongside the entry probability; the equilibrium pins down both. If she did not enter, stage 3 is an ordinary spot stage, exactly like stage 2: the informed trader with probability π , a liquidity trader otherwise.

4.4 Equilibrium

Definition 1 (Equilibrium). An *equilibrium* is a manipulator strategy (her entry probability p_M and her manipulation probability α), a quoting rule for each market maker, and a system of beliefs, taking the informed trader's rule in (3) and liquidity-trader behavior as given, such that:

1. *sequential rationality*—the manipulator's trades maximize her expected profit given the makers' quotes and the beliefs;
2. *zero-profit quoting*—each maker's bid and ask equal the expected value conditional on the order it faces, subject to a non-crossing book (the ask at least the bid), so that where the unconstrained break-even quotes would cross, the maker quotes a zero spread and earns the resulting rent;
3. *belief consistency*—beliefs are formed from the strategies by Bayes' rule on every information set reached with positive probability.

Where condition (i) leaves more than one manipulation probability α consistent with the quotes, we select the smallest one that is stable to small perturbations of the conjecture, following Zhu (2014).

4.5 Price Paths and Notation

Each spot order has one of three outcomes—a buy $+$, a sell $-$, or no trade $\emptyset$ —and we write $O = \{+, -, \emptyset\}$ for this set. At an ordinary spot stage some order always trades (no trade has probability zero), so the outcome is a buy or a sell; write $O_2 = \{+, -\}$ for this smaller set. A spot price path h records the realized outcomes—the stage-2 outcome h_2 , and after the settlement window the pair h_2h_3 . Given a price path h , the posterior probability of the up state is $\rho_h = \mathbb{P}(V = +\sigma \mid h)$; because V takes only the two values $\pm\sigma$, this one number fixes the posterior mean and variance,

$$\bar{V}_h = \sigma(2\rho_h - 1), \quad v_h = \sigma^2(1 - (2\rho_h - 1)^2) = 4\sigma^2\rho_h(1 - \rho_h).$$

Finally, two price-path probabilities recur. We write q_{h_2} for the probability that the stage-2 price path is h_2 , and $q_{h_2h_3}$ for the *conditional* probability that the settlement window then produces outcome h_3 , given h_2 ; their product $q_{h_2}q_{h_2h_3}$ is the probability of the full price path h_2h_3 .

We restrict attention throughout to the empirically relevant region: adverse selection is moderate, and the settlement noise is small.

Assumption 1. $\varepsilon < \sigma\pi < \frac{1}{2}$.

5 Theoretical Results

This section solves the model and then interprets the equilibrium. We first solve the manipulator's settlement-window push for a fixed entry probability (Section 5.1). We then show that manipulator participation reduces settlement-time spot price discovery (Section 5.2), endogenize entry in the equilibrium and characterize the transfer and liquidity asymmetry it produces (Section 5.3), use a second ordinary spot stage to study the contract horizon as the design lever, and generalize the horizon lever to n ordinary spot stages (Section 5.4). Throughout, proofs are deferred to the appendix.

5.1 The Manipulator's Settlement-Window Push

The prediction-market contract settles on the spot price discovered at the close, so the manipulator, who holds the contract, can move what it pays by trading the spot asset.

She acts in the settlement window: stage 3, the spot trade in the final moments before the close. She reaches the window holding the prediction-market contract she bought at stage 1, and we condition throughout on her having bought it (a long position); the short-contract branch is symmetric. Her one decision is whether to *push*—to buy the spot so as to lift the close toward the side she holds—and how often. The push has both a payoff and a price: it lifts the settlement-time spot price, and with it her contract's expected payoff, but she must pay the spot maker's ask, which

exceeds the price the spot has discovered. Figure 3 lays out the resulting continuation game: the stage-2 spot order leaves an up or a down public posterior, after which the manipulator, if she entered, chooses whether to push.




Figure 3. The stage-3 continuation game, conditional on a stage-1 manipulator buy on the prediction market. The stage-2 spot order produces an up posterior ($h_2 = +$) or a down posterior ($h_2 = -$), each with probability $\frac{1}{2}$. If the manipulator entered at stage 1 (probability p_M), she then either pushes the spot with the buy probability α_{h_2} derived below or stays out; otherwise the settlement-window order is an ordinary spot order (an informed trader or a liquidity trader), exactly as at stage 2.

The object we solve for is a *continuation equilibrium*: a set of stage-3 spot quotes and a stage-3 strategy for the manipulator that fit together. Three requirements make them fit: the spot maker quotes break-even prices, equal to the asset's expected value given how the manipulator is expected to trade, subject to a non-crossing book; the manipulator trades optimally given those prices; and the spot maker's beliefs follow Bayes' rule. Two probabilities drive the analysis. The first, p_M , is the probability the manipulator *enters* the contract at stage 1. The second, α , is the probability she *pushes* the spot at stage 3 once she has entered. We solve this continuation game with p_M held fixed, then let stage 1 choose p_M in Section 5.3.

We begin with the spot market's ordinary Glosten–Milgrom round at stage 2, whose only job is to fix the beliefs that stage 3 starts from. A single spot order arrives: with probability π the informed trader, who has seen the asset and buys when it is worth more than the ask and sells when it is worth less; otherwise a liquidity trader, who trades at random. The spot maker cannot tell the two apart, so it quotes to break even, setting each quote to the asset's expected value given the order it sees. Only the informed trader's order carries news, so a buy raises the posterior that the asset is high and a sell lowers it, and the break-even quotes are the matching posterior means. Lemma 1 states these quotes

and the posteriors stage 3 inherits, with the derivation in Appendix A.2.

Lemma 1 (Stage-2 quotes and posteriors). *The stage-2 break-even quotes are the ask $A_2^S = \sigma\pi$ and the bid $B_2^S = -\sigma\pi$. Each outcome $h_2 \in \{+, -\}$ occurs with probability $q_{h_2} = \frac{1}{2}$ and leaves the spot posterior*

$$\rho_{h_2} = \frac{1 \pm \pi}{2}, \quad \bar{V}_{h_2} = \pm \sigma\pi, \quad v_{h_2} = \sigma^2(1 - \pi^2),$$

with the plus sign after a buy and the minus sign after a sell. The posterior mean $\bar{V}_{h_2}$ is the price the spot market has discovered, and the residual variance v_{h_2} is the uncertainty about the asset that remains.

Stage 2 leaves the spot maker with the posterior ρ_{h_2} from Lemma 1, so stage 3 opens from that belief. The manipulator's stage-3 action remains to be pinned down. She would never sell the spot, because a sell moves the close downward, the wrong way for the long contract she holds, and so can only cost her. She is therefore left to buy, to stay out, or to randomize between the two, so her whole stage-3 strategy collapses to a single number, the probability α_{h_2} that she pushes (buys the spot) after stage-2 outcome h_2 .

Before pricing the interior case, two values of this push probability can be set aside. Take $p_M = 1$ first: the lone stage-3 trader is then the manipulator, who is uninformed about the asset, so her buy reveals nothing and the break-even ask equals the prior mean $\bar{V}_{h_2}$, so her push probability does not affect the price. Take $p_M = 0$ next: she is absent from stage 3 altogether, so there is no push. We then fix $p_M \in (0, 1)$ for the rest of the analysis.

Consider next the price the spot maker sets for the single spot order that arrives, an order whose sender it cannot see. A buy can come from any of three types. It can come from the informed trader, who buys only when the asset is truly high; from a liquidity trader, who buys at random; or from the manipulator, who pushes with probability α_{h_2} to lift the close. Of the three, only the informed trader's buy carries information about the asset, because the liquidity trader and the manipulator both trade for reasons unrelated to value and are therefore uninformed. The spot maker cannot tell which type sent the order, so it prices against this pooled flow and sets each quote to break even,

earning zero in expectation given the order it observes. Imposing that break-even condition pins the quotes down, as Proposition 1 records.

Proposition 1 (Stage-3 quotes and posteriors). *Fix the stage-2 outcome $h_2 \in \{+, -\}$, an interior entry probability $p_M \in (0, 1)$, and a conjectured stage-3 push probability $\alpha_{h_2} \in [0, 1]$. The unique break-even stage-3 quotes are*

$$A_+(\alpha_+; p_M) = \frac{\sigma\pi [p_M\alpha_+ + (1 - p_M)]}{p_M\alpha_+ + (1 - p_M) \frac{1+\pi^2}{2}}, \quad A_-(\alpha_-; p_M) = \frac{-\sigma\pi p_M\alpha_-}{p_M\alpha_- + (1 - p_M) \frac{1-\pi^2}{2}},$$

$$B_+ = 0, \quad B_- = \frac{-2\sigma\pi}{1 + \pi^2}.$$

The break-even ask balances the spot maker's expected gain when it sells to an uninformed buyer (a liquidity trader or the manipulator) against its expected loss when it sells to the informed trader. The corollary below traces how the ask responds as the manipulator is expected to push more aggressively.

Corollary 1 (Stage-3 quote monotonicity). *The ask falls as the manipulator is expected to push harder: for each h_2 , the ask $A_{h_2}(\alpha; p_M)$ is strictly decreasing in the conjectured push α and weakly decreasing in the entry probability p_M (strictly when $\alpha > 0$), while the bid B_{h_2} and the no-trade posterior $\bar{V}_{h_2\emptyset} = \bar{V}_{h_2}$ are unaffected.*

The more often the manipulator is expected to push, the more uninformed buys are pooled in with the informed trader's, so any given buy is less likely to be hers. A buy therefore carries less adverse selection risk, and so the spot maker lowers its ask.

The question now is how hard the manipulator actually chooses to push. She pushes only when pushing pays, so we compare the payoff from buying with the payoff from staying out. The difference is the best-response gap $\Gamma_{h_2}(\alpha_{h_2})$, the rise in her contract's expected payoff from a higher

close net of the spot half-spread $A_{h_2} - \bar{V}_{h_2}$ she pays to push:

$$\begin{aligned}
 \Gamma_{h_2}(\alpha_{h_2}) &= \underbrace{H\left(\frac{A_{h_2}(\alpha_{h_2})}{\varepsilon}\right) - [A_{h_2}(\alpha_{h_2}) - \bar{V}_{h_2}]}_{\text{expected payoff from buying}} - \underbrace{H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right)}_{\text{expected payoff from staying out}} \\
 &= \underbrace{H\left(\frac{A_{h_2}(\alpha_{h_2})}{\varepsilon}\right) - H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right)}_{\text{increase in expected contract payoff from buying}} - \underbrace{[A_{h_2}(\alpha_{h_2}) - \bar{V}_{h_2}]}_{\text{spot cost}}.
 \end{aligned}$$

Note that the price the manipulator pays at stage-1 to enter the prediction market is sunk for the purpose of solving her stage-3 optimality, so it does not enter the gap.

In equilibrium the spot maker's conjecture must be borne out: the push the manipulator actually chooses must equal the one the spot maker priced against. Her α_{h_2} is therefore a *fixed point* of the gap. She stays out when pushing never pays ($\Gamma_{h_2}(0) \leq 0$), pushes for sure when it always pays ($\Gamma_{h_2}(1) \geq 0$), and mixes only at an interior probability where the two exactly tie ($\Gamma_{h_2}(\alpha_{h_2}) = 0$). Because both the gain and the cost shrink as the conjecture rises, the gap can cross zero more than once; following Zhu (2014), the equilibrium keeps the smallest *stable* fixed point, and we write $\mathcal{E}_{h_2}$ for the set of stable fixed points.

Proposition 2 (Selected stage-3 continuation equilibrium). *Fix an interior entry probability $p_M \in (0, 1)$. The set of stable fixed points $\mathcal{E}_{h_2}$ is nonempty for each h_2 , so the selected push $\alpha_{h_2}^* = \min \mathcal{E}_{h_2}$ is well defined.*

Existence is the standard fixed-point argument applied to the continuous gap Γ_{h_2} , and the equilibrium takes its smallest stable crossing. This pins down the continuation $\alpha_{h_2}^*$ for every entry probability p_M , which stage 1 then chooses in Section 5.3.

5.2 Price Discovery Falls

The manipulator's mere presence makes the settlement-time spot price strictly less informative, for any entry probability $p_M \in (0, 1)$.

To make “less informative” precise, we measure how much uncertainty about the asset’s value V the spot price still leaves once stage 3 closes. This is the *expected residual variance* $\mathbb{E}[v_{h_2h_3}]$, the variance of V that remains after conditioning on the discovered price, averaged over outcomes:

$$\mathbb{E}[v_{h_2h_3}] = \sum_{h_2} q_{h_2} \sum_{h_3} q_{h_2h_3} v_{h_2h_3}.$$

By the law of total variance, this residual variance and the variance of the discovered price sum to the asset’s prior variance, $v_0 = \mathbb{E}[v_{h_2h_3}] + \text{Var}(\bar{V}_{h_2h_3})$, with the total fixed at $v_0 = \sigma^2$. A price that tracks V closely leaves little residual variance, so a *higher* residual variance means a *less* informative price. The proposition below shows that any interior entry raises it above the no-entry benchmark $p_M = 0$.¹¹

Proposition 3 (Manipulator participation strictly harms price discovery). *With informed traders present ($\pi > 0$), for every interior entry probability $p_M \in (0, 1)$ and every conjectured profile of stage-3 push probabilities $(\alpha_{h_2})_{h_2 \in O_2}$,*

$$\mathbb{E}[v_{h_2h_3}](p_M) > \mathbb{E}[v_{h_2h_3}](0).$$

The loss runs through two channels. The first is displacement: the manipulator occupies the stage-3 slot with probability p_M , thinning the informed order flow by the factor $1 - p_M$, and this alone makes the loss strict even when she never pushes ($\alpha_{h_2} = 0$). The second is dilution: once she does push ($\alpha_{h_2} > 0$), her uninformed buys pool with the informed trader’s, so a buy reveals less about the asset. The benchmark $p_M = 0$ shuts down both.

This is the mirror image of Zhu’s dark-pool result. In Zhu (2014), the second venue draws relatively uninformed flow away from the exchange, concentrating informed traders there and sharpening the exchange price. Here the linked contract feeds uninformed, settlement-motivated flow into the price-setting market, so the settlement-time spot price becomes less informative. A

¹¹From the spot market’s perspective, setting $p_M = 0$ is equivalent to shutting down the prediction market, since it is the only interaction channel between the two.

derivative that settles on a market's price can therefore degrade the informativeness of the price it references.

5.3 The Manipulation Equilibrium

We then solve the stage-1 entry-and-pricing game. The manipulator chooses how often to enter the prediction market. The prediction-market maker, unable to tell the manipulator from an ordinary buyer, quotes a price. Two results emerge. The first is the distinctive feature of the model. The prediction market posts a *zero* bid-ask spread, yet the prediction-market maker still earns a positive rent while the manipulator earns a positive profit. The second is the welfare counterpart. A single uninformed outsider, the liquidity trader, funds both the rent and the profit.

To solve the game we step back to stage 1. Two kinds of buyer arrive there, and the prediction-market maker cannot tell them apart. One is the manipulator. When ordinary flow has gone against the side she holds, she pushes the spot at stage 3, so her contract is worth a high expected payoff. The other is an uninformed liquidity trader. She has no such push, so her contract may be worth far less. The prediction-market maker prices a single contract knowing it may be either type, and we work out three stage-1 objects in turn: the two types' expected payoffs, the break-even price the prediction-market maker posts, and the manipulator's entry decision.

Start with the prediction-market maker's problem. Fix the manipulator's entry probability $p_M \in (0, 1)$. Both payoffs are built from the contract's settlement value H , defined in (2). There, $H(x)$ is the contract's expected payoff when the posterior mean sits x settlement-noise units above the strike, that is, the probability the contract settles in-the-money. If the manipulator enters, her expected PM contract payoff after stage-2 outcome h_2 is

$$\Pi_M^+(p_M, h_2) = \alpha_{h_2}^* H\left(\frac{A_{h_2}^*}{\varepsilon}\right) + [1 - \alpha_{h_2}^*] H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right).$$

This averages her contract value across two events. When the stage-2 price path has gone against her, she pushes the spot at stage 3. When ordinary flow already carries the price her way, she does

not. The liquidity trader has no such ability, so she faces ordinary settlement, and her expected PM contract payoff after h_2 is

$$\Pi_L^+(p_M, h_2) = \left[\frac{1-\pi}{2} + \pi \rho_{h_2} \right] H\left(\frac{A_{h_2}^*}{\varepsilon}\right) + \left[\frac{1-\pi}{2} + \pi(1-\rho_{h_2}) \right] H\left(\frac{B_{h_2}}{\varepsilon}\right),$$

where ρ_{h_2} is the stage-2 posterior that settlement favors her side. Averaging over the stage-2 outcome h_2 using the outcome probabilities q_{h_2} as weights gives the unconditional expected payoff for each stage-1 type:

$$\Pi_M^+(p_M) = \sum_{h_2} q_{h_2} \Pi_M^+(p_M, h_2), \quad \Pi_L^+(p_M) = \sum_{h_2} q_{h_2} \Pi_L^+(p_M, h_2).$$

The competitive prediction-market maker sees only that someone wishes to buy, so it posts a single price that breaks even over the mix of the two types. That break-even ask is the weighted average of the two payoffs:

$$A_1^{P, BE}(p_M) = p_M \Pi_M^+(p_M) + (1 - p_M) \Pi_L^+(p_M).$$

This break-even ask can fall strictly below $1/2$, which implies a break-even bid strictly above $1/2$ and therefore strictly above the break-even ask. But the prediction market cannot post a crossed book, so the posted ask is floored at $1/2$:

$$A_1^{P*}(p_M) = \max\left\{A_1^{P, BE}(p_M), \frac{1}{2}\right\}, \quad B_1^{P*}(p_M) = 1 - A_1^{P*}(p_M).$$

The floor is a general property of the book, not a restriction we impose. When the break-even already exceeds $1/2$, the floor is slack, the posted ask equals the break-even ask, and the spread is positive, just as in an ordinary adverse-selection market. The striking case is the opposite one, and it is the one the model delivers: when the break-even sits below $1/2$, the non-crossing constraint binds, the prediction-market maker is pinned at $1/2$, and it keeps the gap $\frac{1}{2} - A_1^{P, BE}$ as a strictly

positive rent on every contract sold, even though the posted spread is exactly zero.

Now turn to the manipulator's problem. Her profit from entering, paying the posted ask, and pushing at settlement is her expected contract payoff net of the spot-push cost and the price she paid:

$$S_M^+(p_M) = [\Pi_M^+(p_M) - K^+(p_M)] - A_1^{P*}(p_M),$$

where K^+ is the expected resources she spends moving the spot, summed over the two ordinary stage-2 outcomes $h_2 \in O_2 = \{+, -\}$:

$$K^+(p_M) = \sum_{h_2 \in O_2} \underbrace{q_{h_2} \alpha_{h_2}^*}_{\text{probability of manipulation}} \underbrace{[A_{h_2}^* - \bar{V}_{h_2}]}_{\text{spot manipulation cost}} \geq 0.$$

She chooses her entry frequency p_M to maximize her ex-ante expected profit, the entry probability times the profit per entry:¹²

$$\bar{S}_M^+(p_M) = p_M S_M^+(p_M), \quad p_M^* \in \arg \max_{p_M \in [0,1]} \bar{S}_M^+(p_M).$$

For compactness, write the pure-push cutoff derived in Appendix A.5.2 as

$$p_0 := \left(1 + \frac{2(\sigma\pi - \varepsilon)}{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)} \right)^{-1}.$$

The following proposition characterizes the full equilibrium.

Proposition 4 (The manipulation equilibrium). *Fix $\sigma > 0$, $\pi \in (0, 1)$, and $\varepsilon > 0$ satisfying Assumption 1. The equilibrium of Definition 1 has:*

1. *optimal entry probability*

$$p_M^* = \sqrt{\left(\frac{1 - \pi^2}{1 + \pi^2} \right)^2 + \frac{1 - \pi^2}{1 + \pi^2} p_0} - \frac{1 - \pi^2}{1 + \pi^2} \in (0, p_0);$$

¹²A maximizer exists whenever $\bar{S}_M^+$ is continuous in p_M , which holds below, where $\bar{S}_M^+$ is available in closed form. When the maximizer is not unique, we select the smallest one.

2. *push policy*

$$\alpha_+^* = 0, \quad \alpha_-^* = 1;$$

3. *strictly positive expected profit*


$$\bar{S}_M^+(p_M^*) = \frac{(1 - 2\sigma\pi)p_M^{*2}}{4p_0} > 0;$$

4. *zero-spread quotes for the prediction market:*


$$A_1^{P*} = B_1^{P*} = \frac{1}{2}.$$

The manipulator enters with an interior probability p_M^* . She enters often enough to profit, but not so often that a worse prediction-market ask and a smaller settlement gain from her stage-3 push erode her profit. Her push policy sits at a corner. She never pushes after a stage-2 buy ($\alpha_+^* = 0$), because ordinary flow has already carried the price her way, and there is nothing to manipulate. She always pushes after a stage-2 sell ($\alpha_-^* = 1$), because flow has gone against her side and a push can still move the settlement her way.

If both the prediction-market maker and the manipulator profit at a zero spread, a third party must fund them. That party is the uninformed liquidity trader. She buys at the same $1/2$ but has no settlement-window push of her own, so she expects to receive less than what she pays. Why does she lose in expectation, holding no information and no view? The pooling that does it sits at the *spot* maker, not the prediction-market maker. The manipulator is long the “up” contract and would lift the settlement price by buying the spot. A settlement-window buy becomes less informative, so by Corollary 1 the spot maker lowers its ask while the bid is unchanged. The result is a settlement bias *against* the “up” holder, including the liquidity trader even though she is innocent. She need not meet the manipulator, hold any view about the spot asset, or even know the spot market exists; simply holding the “up” contract exposes her to the spot-maker quote response. The harm is cross-market (Figure 4): a prediction-market position triggers a spot-maker quote response that



(a) The zero-spread rent: the break-even price lies below $1/2$, so the non-crossing floor pins the posted price at $1/2$ and the prediction-market maker keeps the gap as a positive rent.



(b) The cross-market channel: a prediction-market “up” buy makes a settlement-window spot buy look possibly manipulative, so the spot maker lowers its ask; the settlement is biased against the “up” side (the holder collects $< 1/2$) while the spot ask tightens (more liquid).

Figure 4. The zero-spread rent and the cross-market channel. Panel (a) shows where the prediction-market maker’s rent comes from; panel (b) shows how the uninformed liquidity trader’s loss is inflicted through the spot maker’s quotes.

tilts the settlement against that position, which harms the contract holder but makes the spot market more liquid.

The following proposition gives the accounting.

Proposition 5 (Liquidity-trader loss decomposition). *At the equilibrium of Proposition 4, evaluated at the interior entry probability $p_M^* \in (0, p_0)$, the liquidity trader’s ex-ante loss splits exactly into*

three components:

$$(1 - p_M^*) L^+(p_M^*) = R_P^{+*} + p_M^* S_M^+(p_M^*) + p_M^* K^+(p_M^*),$$

where the prediction-market maker's rent, the liquidity trader's per-order loss, and the spot push cost are

$$R_P^{+*} := A_1^{P*} - A_1^{P, BE}(p_M^*) = \frac{p_M^* (\sigma\pi - \varepsilon)}{4\varepsilon} > 0,$$

$$L^+(p_M^*) := \frac{1}{2} - \Pi_L^+(p_M^*) = \frac{(1 - \pi^2) \sigma\pi p_M^*}{4\varepsilon [(1 - \pi^2) + (1 + \pi^2) p_M^*]} > 0,$$

$$K^+(p_M^*) = \frac{1}{2} \frac{\sigma\pi (1 - \pi^2) (1 - p_M^*)}{(1 - \pi^2) + (1 + \pi^2) p_M^*}.$$

The identity shows that the liquidity trader's expected loss funds the two rents and one cost on the right-hand side. The first rent is the prediction-market maker's R_P^{+*} , the gap between the posted $1/2$ and the break-even ask that the prediction-market maker keeps because it cannot quote any tighter under the non-crossing constraint. The second is the manipulator's ex-ante profit $p_M^* S_M^+$ while paying the same $1/2$. The third term, $p_M^* K^+$, is the manipulator's ex-ante cost of moving the spot in the settlement-window push, which helps the spot maker finance the adverse-selection loss to the informed trader.

Having pinned down who pays, we next ask when manipulation incentives are strongest.

Corollary 2 (Comparative statics). *The signs of the comparative statics are summarized in the table below.*

<i>Quantity</i>	<i>Symbol</i>	$\partial/\partial\varepsilon$	$\partial/\partial\sigma$	$\partial/\partial\pi$
<i>Pure-push cutoff</i>	p_0	> 0	< 0	< 0
<i>Optimal entry probability</i>	p_M^*	> 0	< 0	< 0
<i>Stage-3 ask</i>	$A_-^*(p_M^*)$	< 0	> 0	> 0
<i>Spot manipulation cost at p_M^*</i>	$A_-^*(p_M^*) - \bar{V}_-$	< 0	> 0	> 0
<i>Ex-ante manipulation profit</i>	$\bar{S}_M^+(p_M^*)$	> 0	< 0	< 0

When settlement noise ε rises, the adverse outcome is easier to rescue, so the manipulator enters more often and p_M^* rises. Because she now enters more often, a stage-3 buy after a stage-2 sell is more likely to be manipulative and therefore less informative about the fundamental, so the stage-3 ask and the spot-push cost both fall. Even so, expected profit rises, because the gain from entering more often outweighs the lower profit per entry.

When fundamental volatility σ or information frequency π rises, the adverse outcome is harder to rescue. Each moves the stage-2 posterior further from the strike, so the settlement-window push is more costly. The manipulator enters less often, and expected profit falls.

The full comparative-statics derivations are in Appendix A.5.3.

5.4 The Horizon Lever

A venue cannot change a trader's signal or a manipulator's budget, but it can change one thing directly: the contract's *horizon*, the span of trading before the price settles. The horizon can be a design lever against manipulation. Each ordinary trading round before settlement lets the price absorb more information, so the posterior entering the settlement window sits further from the strike and a fixed-size push stays pivotal on a smaller set of potential paths. As the horizon grows, manipulation grows *rarer* and the settlement price *more* informative; lengthen the horizon without

bound and manipulation vanishes outright. We first solve the smallest extension, one extra round, where the mechanism is cleanest, then the general case of n ordinary rounds.

We lengthen the contract's horizon by inserting $n \geq 1$ ordinary trading rounds between the manipulator's entry and the settlement window. Each round is an ordinary Glosten–Milgrom round: an informed trader arrives with probability π and trades on her signal as in (3), and otherwise a liquidity trader buys or sells at random, so each round adds one buy + or one sell – to the public record. The manipulator still trades only in the settlement window. The baseline is the one-round case, $n = 1$.

A stage-2 price path with n rounds is denoted by a string of n buys or sells, $h_2 = h_{2,1} \cdots h_{2,n}$. What the market learns from it is summarized by a single number, the *order imbalance* $d(h_2)$, the number of buys minus the number of sells. The imbalance is a sufficient statistic for the posterior entering settlement: two price paths with the same imbalance leave the market at the same posterior, so a stage-2 price path can be summarized by d alone.

A fixed-size push moves the settlement price by a bounded amount, so it can change which side wins only when the price already sits near the strike—that is, only when the imbalance is near balance. Once the order flow has moved the imbalance far to one side, the outcome is settled and the manipulator stays out. The horizon works by making a near-balance imbalance rarer. Figure 5 lays out the timeline.




Figure 5. The extended game with n ordinary spot rounds between entry and settlement. The manipulator decides whether to enter at stage 1, n ordinary Glosten–Milgrom rounds then reveal order flow, and at stage 3 the manipulator (if she entered) pushes the spot price before the contract settles. The baseline (Section 5.3) is the case $n = 1$.

We solve the smallest case first—two rounds, where the few price paths can be enumerated and the mechanism is clearest—then the general n -round case. With two ordinary rounds the order flow can come out four ways. When the two orders agree—both buys or both sells—they move the

posterior far enough from the strike that the settlement is already decided. When they disagree—a buy and a sell—they cancel and leave the posterior at the strike, where a push can still change the settlement. So manipulation survives only on the two disagreeing paths. Figure 6 traces the four paths, and Appendix A.6.1 records the posteriors they reach.




Figure 6. The lengthened game with two ordinary spot rounds before the settlement window, conditional on a stage-1 buy. The two ordinary orders (stage 2.1 then stage 2.2) produce four stage-2 price paths. The agreeing price paths ++ and -- move the posterior far enough from the strike that a push cannot change the settlement, so the manipulator stays out; the disagreeing price paths +- and -+ leave it at the strike, where the outcome stays undecided and she pushes.

The equilibrium of the two-round game is as follows.

Proposition 6 (Equilibrium with two ordinary spot rounds). *Fix $\sigma > 0$, $\pi \in (0, 1)$, and $\varepsilon > 0$ satisfying Assumption 1. The two-round game has a unique equilibrium (Definition 1), with a unique optimal entry probability $p_M^{*(2)} \in (0, 1)$. In this equilibrium:*

1. *she pushes only after the two ordinary orders disagree, and never after they agree:*

$$\alpha_{++}^* = \alpha_{--}^* = 0, \quad \alpha_{+-}^* = \alpha_{-+}^* = 1;$$

2. *price discovery strictly improves relative to the baseline:*

$$\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (p_M^{*(2)}) < \mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (p_M^{*(1)}).$$

The two-round case isolates the horizon lever in miniature: the extra ordinary round filters out the agreeing price paths, leaving only the disagreeing price paths vulnerable to manipulation. We

now solve the same logic for general n .

Proposition 7 (Equilibrium with $n > 2$ ordinary spot rounds). *Fix $n > 2$, $\sigma > 0$, $\pi \in (0, 1)$, and $\varepsilon > 0$ satisfying Assumption 1. The n -round game has a unique equilibrium (Definition 1), with a unique optimal entry probability $p_M^{*(n)} \in (0, 1)$. In this equilibrium:*

1. *she pushes only when the order imbalance is near balance, leaving the price close to the strike:*

$$\alpha_{h_2}^{*(n)} = \begin{cases} 1, & n \text{ is odd and } d(h_2) = -1, \\ 1, & n \text{ is even and } d(h_2) = 0, \\ 0, & \text{otherwise;} \end{cases}$$

2. *price discovery strictly improves with every added ordinary spot stage:*


$$\mathbb{E} \left[v_{h_2 h_3}^{(n)} \right] (p_M^{*(n)}) < \mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (p_M^{*(n-1)}).$$

3. *the entry probability differs by the parity of n :*


$$p_M^{*(n)} = \begin{cases} p_M^{*(1)}, & n \text{ is odd,} \\ p_M^{*(2)}, & n \text{ is even,} \end{cases}$$

Two things matter. First, however long the horizon, a push pays on the price paths that leave the price closest to the strike but still on the wrong side of it: the balanced ($d = 0$) price paths when n is even and the one-step-below ($d = -1$) price paths when n is odd. Second, price discovery improves with every round: each additional round gives the market one more reading of the order flow, so the settlement price tracks the asset's value more closely, and the expected residual variance $\mathbb{E}[v_{h_2 h_3}^{(n)}]$ introduced in Section 5.2 falls. The entry probability itself does not settle down—it oscillates with the parity of the horizon, but that oscillation does not govern the contract's safety. What governs safety is how often a push can be pivotal at all.

That a push pays only on near-balanced price paths is why manipulation disappears as the horizon grows. The imbalance follows a random walk, moving ± 1 each round (Figure 7). A push is pivotal only if the imbalance ends within one step of balance, but its dispersion grows with the horizon—the unconditional standard deviation of the imbalance is of order n —while the width of the pivotal band stays fixed. The probability of ending inside the band therefore falls to zero, and the manipulator’s expected profit with it.



short horizon: the walk ends in the band ($|d| \leq 1$)—a push can still change the settlement



long horizon: the walk is likely to end far from balance, outside the band, so a push no longer affects the settlement

Figure 7. The pre-settlement order imbalance d as a ± 1 random walk over the ordinary stages. The shaded strip is the pivotal band: a settlement-window push can change the settlement only when the walk ends within one step of balance ($|d| \leq 1$)—exactly $d = 0$ when n is even and $d = -1$ when n is odd. As the horizon n grows the walk spreads out (the imbalance has standard deviation of order n) while the band stays fixed, so the chance of ending inside it—the pivotal probability $m^{(n)}$ —falls to zero.

Proposition 8 (Manipulation vanishes as the horizon grows). *As the horizon grows, manipulation becomes both rarer and worthless. Write $m^{(n)}$ for the chance a push is still pivotal, $S_M^{+(n)}$ for the manipulator’s profit conditional on entry, and $\bar{S}_M^{+(n)}$ for her unconditional expected profit at the optimal entry. All three vanish:*

$$\lim_{n \uparrow \infty} m^{(n)} = 0, \quad \lim_{n \uparrow \infty} S_M^{+(n)}(p_M^{*(n)}) = 0, \quad \lim_{n \uparrow \infty} \bar{S}_M^{+(n)}(p_M^{*(n)}) = 0.$$

The horizon is therefore a design lever for market safety. Lengthening a contract’s horizon does not make manipulation impossible—a large enough push always moves the settlement price—but it makes a fixed-size push almost never pivotal. A contract is safe not when a manipulator cannot act,

but when enough price discovery runs before settlement that a single trader is unlikely to affect the outcome.

A robust settlement reference is a second, equally powerful lever. Our contract settles on the spot price at a single instant—the close—which is exactly what a final-seconds push can move. A reference that aggregates instead—averaged over a window, or struck in a closing auction that pools many participants’ orders—forces the manipulator to move the whole aggregate rather than one print; a time-weighted oracle, for one, raises the manipulation cost roughly in proportion to the averaging window (Aspembitova and Bentley, 2022). The two levers address manipulation from opposite sides: a longer horizon adds price discovery *before* settlement, while an aggregating reference spreads the settlement price *across* orders or time. The binary Nasdaq-100 index options the SEC approved in April 2026 (SR-MRX-2026-05) sit on the safe side of both: their expirations run from weekly to multi-year, and they settle on the index value struck in the Nasdaq Closing Cross, the exchange’s closing auction, which pools the day’s closing interest into one print and is correspondingly hard to push.¹³ Our five-minute, single-instant contract sits at the opposite corner.

6 Empirics

This section takes the model’s predictions to the data. We first state what settlement manipulation should leave in the data according to the theory (Section 6.1). We then describe the sample (Section 6.2) and present the empirical evidence in two parts: Section 6.3 documents a near-settlement *push* in the spot market, defined as a burst of one-sided trading in the final seconds before a contract settles, and shows it bears the predicted signature: it appears only for the five-minute contract, concentrates in the cycles the prediction market has left undecided, causes price reversals within seconds of settlement, and improves spot liquidity; Section 6.4 then shows the push is

¹³Nasdaq MRX, LLC, *Order Granting Accelerated Approval of a Proposed Rule Change to Adopt New Options Rule 3B to List and Trade Binary Broad-Based Index Options*, Exchange Act Release No. 34-105342, File No. SR-MRX-2026-05 (Apr. 30, 2026). The settlement value is the Nasdaq-100 Index as reported at the conclusion of the Nasdaq Closing Cross (Nasdaq Equity 4, Rule 4757); the contracts are cash-settled, European-style binary options listed in weekly, monthly, quarterly, end-of-month, and long-term (12–60 month) expirations.

deliberate manipulation: it overturns resolutions even in near-certain cycles, transfers wealth from retail holders to a small cohort of manipulators, and cannot be explained solely by prediction-market makers' hedging.

6.1 Model Predictions

The first set of predictions concerns the footprint manipulation leaves in the market. If settlement manipulation exists, it must appear as a directional push in the final seconds before the close: a burst of one-sided order flow on the spot venue the oracle samples, large enough to move the spot price (Proposition 4). Because that flow carries no information about fundamentals, its price impact should be transitory: the move should revert once settlement has passed, because the settlement-window price is less informative (Proposition 3), in contrast to the permanent impact of informed trading.

The same lack of information content shapes liquidity on the spot venue. The manipulator's uninformed order pools with informed flow and lowers the adverse selection any individual trade carries, so spot market makers supply *more* liquidity into the close even as activity surges (Corollary 1). Finally, the entire footprint should be governed by whether a push can change anything: it concentrates in cycles the prediction market has left undecided near the close, and in the shortest-horizon contract, where the asset's own price discovery has had the least time to unfold and a fixed-size push is most likely to be pivotal (Proposition 7 and Proposition 8).

The second prediction concerns who pays. A successful push can overturn the resolution the market had priced, even in near-certain cycles we later identify in the data (Proposition 4). The contract is then mispriced relative to the manipulated resolution, and the loss falls on whoever holds the losing side at the close: liquidity traders, as predicted in the model. Manipulator profits are, to first order, liquidity traders' losses, according to the model's accounting identity (Proposition 5).

6.2 Data

The analysis combines three data sources: on-chain Polymarket trade records for the crypto up/down contracts, off-chain Polymarket order-book snapshots, and trades and quotes for BTC spot on Binance. Binance is the reference spot market throughout: it carries the largest BTC spot trading volume and arguably has the largest impact on the Chainlink reference price against which the Polymarket contracts resolve, so it is the natural venue on which a settlement manipulator would trade. Appendix D validates this choice directly, showing that the Binance midquote tracks the Chainlink resolution price to within a few basis points.

- *Polymarket on-chain trades.* Every order fill executed on Polymarket’s exchange contracts settles on the Polygon blockchain, from which we extract the crypto up/down markets: over 60 million fills on the BTC five-minute contract alone. Each fill carries the block timestamp, the wallet identifiers on both sides, the asset on each side (one side is USDC, the other an Up- or Down-token), and the amounts deposited and received, which together recover the trade’s price (in dollars per share, between 0 and 1), size, direction, and aggressor.¹⁴ We link each fill to its market’s cycle, strike, and realized resolution. Volume is measured two ways: *share volume*, the number of shares changing hands (each share pays \$1 if its side wins, so the share count is the natural unit of exposure), and *dollar volume*, the cash transferred. We sign each fill by the aggressor’s direction, expressed in the Up token (an aggressive Up buy or Down sell is bullish-Up), so net order flow over a window measures directional pressure on the contract.
- *Polymarket order book.* Polymarket’s market-data feed delivers both full limit-order-book snapshots (the prevailing best bid and ask alongside several dozen price levels of resting depth on each side) and single-level updates whenever an existing level’s resting size changes, for

¹⁴The on-chain log produces multiple records for each economic Polymarket trade: one aggregate record summarizing the aggressor’s fill, plus one record for each resting order the aggressor consumes. We keep the single aggregate record per trade, so volume is never double-counted; on that record, the aggressor’s deposited and received assets jointly identify both the contract side traded and the dollars exchanged. For wallet-level position and PnL tracking (Section 6.4.4), we instead sum the resting-side counterparty across both record classes, so each wallet’s net position is recovered exactly once.

both tokens of each market. The feed covers roughly the final seven weeks of the sample. From it we measure the Up-token midquote, the bid–ask spread, and resting depth near the mid.

- *Binance BTC spot trades and order book.* The full tick-by-tick trade records and order-book snapshots for the reference asset. Each trade carries price, size, and aggressor side, from which we build signed spot order flow; each book snapshot carries the five best bid and ask levels with their resting quantities, timestamped at 100-millisecond precision, from which we build time-weighted spreads and depth. Both are aggregated to ten-second buckets aligned with the five-minute contract clock.

Sample window and rollout periods. The sample covers July 1, 2025 through April 8, 2026, starting about three months before the launch of the oracle-resolved Polymarket up/down contracts we study and ending about two months after the five-minute contract launched. We divide the sample into three periods based on the contract rollout schedule: the pre-launch baseline (P1), July 1 through October 8, 2025 (100 days); the intermediate period with the fifteen-minute and four-hour contracts live (P2), October 9, 2025 through February 11, 2026 (126 days); and the post-launch period with the five-minute contract also live (P3), February 12 through April 8, 2026 (56 days).¹⁵

Summary statistics. Table 1 reports summary statistics for the per-cycle trading activity over the five-minute contract’s live window.¹⁶

Two observations stand out. First, the prediction market is small relative to its underlying spot: the median cycle trades about 105 thousand \$1-face shares on Polymarket against \$3.15 million

¹⁵Four narrow windows within the post-launch period require care: in three (late February 2026) the oracle-resolved Polymarket crypto contracts stop accepting trades while the rest of the platform trades normally (consistent with brief oracle-integration interruptions), and in a fourth (March 2026) the platform halts entirely. Cycles in these windows carry no on-chain Polymarket trades, so they drop out of every analysis that requires Polymarket data; the Binance-side panels retain them. On the spot side, we exclude October 10, 2025 (a market-wide flash crash) from the intra-cycle Binance figures, because a transient hundred-dollar spread during that episode would dominate a time-weighted average.

¹⁶Each five-minute contract is listed for trading several hours before its live resolution window opens, but about 93% of five-minute Bitcoin trades fall inside the live window (the 300 seconds between open and close), with the remainder split roughly evenly between pre-window positioning and the brief post-close interval before settlement. All cycle-level aggregates restrict to the live window.

Table 1. Per-cycle summary statistics, BTC five-minute contract. This table reports per-cycle trading activity on the BTC five-minute contract and the underlying Binance spot market, one observation per cycle in the contract's live window (February 12 – April 8, 2026); N is the number of cycles and the remaining columns report the cross-cycle distribution. The *Polymarket* rows aggregate the contract's own trades over the 300-second live window: *Trades* is the per-cycle taker trade count, *Volume* the number of shares traded (each carrying \$1 of face value), and *Distinct traders* the number of unique wallets with a fill. The *Binance* rows aggregate the reference asset's spot tape over the same cycles: *Trades* and *Volume* (dollar volume) are per-cycle sums, and *Absolute return* is the absolute open-to-close change in the log midquote over the five-minute window (basis points).

	N	Mean	SD	Min	P25	Median	P75	Max
<i>Polymarket BTC 5-minute contract</i>								
Trades	16,073	3,912	1,228	2	3,084	3,811	4,642	12,622
Volume (1000 shares)	16,073	108.9	36.4	0.0	84.3	104.9	129.8	323.9
Distinct traders	16,073	1,689	374	4	1,467	1,685	1,918	3,332
<i>Binance BTC spot</i>								
Trades	16,128	15,793	14,039	765	7,131	11,676	19,290	223,213
Volume (\$M)	16,128	5.08	7.01	0.17	1.79	3.15	5.86	284.69
Absolute return (bps)	16,128	10.16	11.81	0.00	2.88	6.58	13.51	376.29

of Binance volume. At first glance this asymmetry cuts against a manipulation story: if the prize pool is small relative to the market one must move, why bother? Two considerations restore the incentive. The gap is a factor of thirty, not several orders of magnitude, so the pool at stake is economically meaningful. More importantly, the five-minute-volume comparison is not the relevant one: the manipulator does not need to match five minutes of spot trading, only to deliver the final push, a burst of one-sided trading in the last seconds that carries the close across the strike. Second, ordinary five-minute spot moves are modest: the median absolute open-to-close return is 6.6 bps, with the inter-quartile range from 2.9 to 13.5 bps. A near-settlement spot push of a few basis points can therefore be large enough to clear the strike in a near-the-money cycle.

6.3 Does manipulation leave a footprint?

Section 6.1 tells us what to look for: if manipulation exists, the close of each cycle should carry a footprint with four parts: a burst of directional *order flow* (the push itself), an outsized *absolute return* (the push moves the price), improved spot *depth* (the push is uninformed, so spot makers supply more liquidity, not less), and a post-settlement *reversal* (for the same reason, the price impact does not last). We look for this footprint in three steps. We first show it emerges with the

launch of the five-minute contract. A launch coincidence alone cannot establish manipulation, so we then turn to the identification: within the post-launch period, the footprint concentrates in the cycles where a push pays off, the cycles still undecided near the close. We close with the model's horizon prediction: the footprint should weaken as the contract horizon lengthens, and indeed the fifteen-minute contract, live earlier and settling on the same rule, exhibits only a strongly attenuated version of it.

6.3.1 The footprint emerges with the five-minute launch

Near-settlement activity spike. We divide each five-minute cycle into 30 ten-second bins, indexed 0 through 29, and compute within each bin three of the footprint's four parts: absolute order flow, absolute return, and depth within \$0.10 of the midquote. For each bin we then average across all cycles in a period and compare the pre-launch baseline (P1) against the five-minute-live period (P3); the intermediate period (P2), in which only the fifteen-minute and four-hour contracts trade, returns when we study the fifteen-minute contract on its own (Section 6.3.3).

Figure 8 plots the resulting intra-cycle profiles. The pattern is striking. In bin 29, the last ten seconds of the cycle, all three metrics spike once the five-minute contract is live (the +5m series); in the pre-launch baseline, bin 29 looks no different from the rest of the cycle (we call the first 25 bins the *body*). The jump is confined to the final bin, so the contrast at the close is not a broad shift in Binance trading.¹⁷

We formalize the visual pattern with a pre/post comparison of the bin-29 activity, comparing the pre-launch baseline (P1) against the period in which the five-minute contract is live (P3):

$$\log y_{c,b} = \gamma_b + \alpha_d + \mu_h + \delta_{P3} \mathbf{1}\{b=29\} \cdot \mathbf{1}\{P3\} + \varepsilon_{c,b},$$

where $\log y_{c,b}$ is the log of the metric in cycle c and bin b , restricted to the body bins and the final-10s bin, $b \in \{0, \dots, 24, 29\}$, so the comparison is clean between the near-settlement activity and the body.

¹⁷There are minor spikes at the minute-ends. They are likely driven by trading algorithms that react to the minute mark.




Figure 8. Binance BTC intra-cycle metrics per 5-min cycle, by period. This figure plots the intra-cycle profile of three Binance BTC metrics, by period. Each panel folds every 10s Binance BTC bucket onto its position in the rolling 5-minute clock cycle and overlays, for each period, the metric pooled over all 5-minute cycles in that period (see Section 6.3.1). *Left:* absolute order flow per cycle (\$M), the per-cycle |net signed taker dollar flow| in the bin. *Center:* absolute 10s return ($|\Delta \log \text{mid}|$, bps). *Right:* time-weighted depth within ± 10 ticks of the mid, cumulative resting size (bid + ask) at price levels within 10 ticks (\$0.10) of the mid. The dotted vertical line marks the final 10s before close. Two periods are overlaid: pre-Polymarket (P1, Jul–Oct 2025) and the five-minute-live period (P3, Feb 12 – Apr 8, 2026); the intermediate period with only the fifteen-minute and four-hour contracts live (P2) is omitted.

α_d and μ_h are date and hour-of-day fixed effects; γ_b is a bin fixed effect that absorbs the pre-existing intra-cycle profile. The specification pools all cycles: we do not differentiate overlapping from non-overlapping closes here; the overlap split returns in the fifteen-minute analysis (Section 6.3.3) and, interacted with moneyness, in the within-P3 refinements of Appendix C. The coefficient δ_{P3} is the average bin-29-versus-body activity increase once the five-minute contract is live, relative to the pre-launch baseline.

Table 2 reports the estimates. Once the five-minute contract is live, the typical cycle’s bin-29 absolute order flow is 50% higher relative to the body than before the launch ($\hat{\delta}_{P3} = 0.406$ log points), its absolute return is 15% higher (0.140), and its near-mid depth 5% higher (0.048), all significant at the 1% level. Three parts of the footprint thus appear at the close exactly when the five-minute contract begins to settle there.

Table 2. Near-settlement activity: pre-Polymarket vs five-minute period. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and to the pre-Polymarket period (P1, Jul–Oct 2025) and the five-minute-live period (P3, Feb 12 – Apr 8, 2026); the intermediate period in which only the fifteen-minute and four-hour contracts are live (P2) is excluded. The dependent variable is $\log y_{c,b}$, the log of the metric named in the column header: *Order flow*, per-cycle absolute net signed taker dollar flow ($|\sum_i d_i p_i a_i|$, $d = +1$ buy / -1 sell); *Abs. return*, absolute 10 s log return ($|\Delta \log \text{mid}|$, bps); *Depth*, time-weighted resting size within ± 10 ticks of the mid. δ_{P3} is the coefficient on $\mathbf{1}\{b=29\} \cdot \mathbf{1}\{P3\}$: the bin-29-versus-body activity increase once the five-minute contract is live, relative to the pre-Polymarket baseline. Fixed effects are date, hour-of-day, and bin; the specification pools all cycles and does not differentiate overlap from non-overlap closes. Coefficients are log differences. One observation per (cycle, bin) cell. HC1 standard errors in parentheses are clustered at the cycle level. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Order flow	(2) Abs. return	(3) Depth
δ_{P3}	0.406*** (0.019)	0.140*** (0.018)	0.048*** (0.003)
Date FE	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes
Observations	1,168,118	738,313	1,168,114
Clusters (cycle)	44,928	44,696	44,928

Post-settlement reversal. The fourth part of the footprint is the reversal: manipulation flow is uninformative, so its price impact should not last. We test this by regressing each lagged ten-second return around the settlement on the final-10s return r_{29} at the cycle level.

Separately at each lag j , we estimate

$$r_{j,c} = \alpha_d + \mu_h + \beta r_{29,c} + \delta_{P3} r_{29,c} \mathbf{1}\{P3\} + \varepsilon_{j,c},$$

on the pre-Polymarket (P1) and five-minute-live (P3) cycles. The dependent variable $r_{j,c}$ is the lag- j ten-second return of cycle c measured relative to its close: positive lags map to bins 0–5 of the next cycle (post-settlement), negative lags to bins 25–28 of the same cycle (the run-up into the close), and $j = 0$ is the bin-29 reference itself, omitted because its slope is 1 trivially. The regressor $r_{29,c}$ is the final-10s (bin-29) return. β is the pre-Polymarket reversal slope and δ_{P3} the change in that slope once the five-minute contract is live; α_d and μ_h are date and hour-of-day fixed effects, and standard errors are clustered by date. A reverting price is a negative slope; the key lag is $j = +1$, the first ten seconds after the close, where a transitory manipulation push predicts reversal.

Figure 9 overlays the reversal slope γ_j for the two periods. A sharp reversal appears in the first




Figure 9. Post-settlement reversal slope γ_j : pre-Polymarket vs five-minute period. This figure plots, at each lag j , the slope γ_j from a per-period pooled OLS of the lag- j 10s return on the final-10s return r_{29} (date and hour-of-day fixed effects, standard errors clustered by date), overlaid for the pre-Polymarket period (P1) and the five-minute-live period (P3); the intermediate +15m/4h period (P2) is omitted. Markers are γ_j ; error bars are $\pm 1.96 \times$ cluster-robust SE; $j = 0$ is omitted. The dotted vertical marks the close. Negative j is bin $29 + j$ of the same cycle; positive j is bin $j - 1$ of the next cycle.

ten seconds after the close, and only with the five-minute contract: γ_{+1} is strongly negative once the contract is live and indistinguishable from zero in the pre-launch baseline, while at $j \geq +2$ the slopes are flat in both periods.

Table 3 confirms it. Against a pre-launch baseline with no reversal at the close ($\beta = 0.05$ at $j = +1$, indistinguishable from a random walk), the five-minute contract shifts the reversal slope by $\delta_{P3} = -0.15$, so the post-close slope under the five-minute contract is $\beta + \delta_{P3} = -0.10$ at $j = +1$: about a tenth of the final-10s spot move reverses within the next ten seconds. The shift is specific to the first post-close lag: δ_{P3} is flat at $j \geq +2$. Taken together, all four parts of the footprint appear with the five-minute contract, and only then. A before/after comparison, however, cannot on its own distinguish manipulation from any other change in close-time trading that coincides with the launch. The identification comes from the cross-section within the post-launch period, to which we now turn.

Table 3. Post-settlement reversal: pre-Polymarket vs five-minute period. This table reports the cycle-level OLS $r_j = \alpha_d + \mu_h + \beta r_{29} + \delta_{P3} r_{29} \mathbf{1}\{P3\} + \varepsilon$, estimated separately at each lag $j \in \{-3, -2, -1, +1, +2, +3\}$, on the pre-Polymarket period (P1, Jul 1 – Oct 8, 2025) and the five-minute-live period (P3, Feb 12 – Apr 8, 2026); the intermediate period in which only the fifteen-minute and four-hour contracts are live (P2) is excluded. β is the pre-Polymarket reversal slope on r_{29} and δ_{P3} its change once the five-minute contract is live (a negative slope is reversal). Date and hour-of-day fixed effects; HC1 standard errors in parentheses are clustered by date. Negative j = same-cycle bin ($29 + j$); positive j = bin ($j - 1$) of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
β	-0.014 (0.034)	0.015 (0.014)	0.051*** (0.018)	0.049 (0.032)	-0.033* (0.019)	-0.002 (0.019)
δ_{P3}	0.015 (0.036)	0.007 (0.019)	-0.012 (0.020)	-0.152*** (0.040)	0.007 (0.026)	0.002 (0.023)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44,772					
Clusters (UTC date)	156					

6.3.2 The footprint concentrates in NTM cycles

The before/after evidence so far shows a four-part footprint that coincides with the five-minute launch, but a launch coincidence cannot by itself separate manipulation from any change in close-time trading. The identification comes from a cross-sectional cut *within* the five-minute period. As the model predicts, if the push is manipulation, all four parts of the footprint should be largest in cycles where the outcome is still undecided near the close — there a small spot push can flip the resolution, so the manipulator's payoff is greatest — and muted in cycles already decided, where a push changes nothing. Because this contrast is taken within a single period, it differences out any time trend common to the launch.

Near-the-money classification. We classify each post-launch five-minute cycle by its prediction-market price near settlement. Specifically, for each cycle, we take the Up-token price 10 seconds before the settlement and label the cycle *near-the-money* (NTM) if the Up-token price is between 0.40 and 0.60 (the contract is still a coin-flip, so a push across the strike can flip the outcome) and *far-from-the-money* (FFM) otherwise. About 6% of cycles are NTM.¹⁸

Near-settlement spike in NTM cycles. Figure 10 plots three intra-cycle metrics — absolute order flow, absolute return, and depth — on both venues, separately for NTM and FFM cycles. The figure carries two messages. First, the Binance spike concentrates in NTM cycles: order flow and absolute return rise at the close in both groups, but markedly more where the contract is still undecided and a push can flip the resolution. Second, liquidity diverges across the venues into the NTM close: Binance makers *add* near-mid depth while Polymarket makers *withdraw* it. Two secondary observations complete the picture. The binary's own price moves more in NTM cycles for a mechanical reason: a price near 0.5 must travel to 0 or 1 as the contract resolves, while a decided contract's price is already there. Its order flow, by contrast, spikes *less*: close-time flow on

¹⁸The ~4% of P3 cycles with no qualifying trade in the 10–12 s window are dropped as unclassified; these are systematically the sparse-trading cycles. The tight window trades coverage for freshness: a wider window would retain them but allow the classifying price to be stale by up to several minutes, mislabeling a cycle when the spot moves in the interim. Pre-launch cycles (P1, P2) have no on-chain 5m trades and are unclassified by construction.

the binary is dominated by the one-sided sweep of an already-decided contract, largest exactly when the outcome is no longer in play, and in contested cycles the directional flow appears on Binance instead.


To quantify the NTM-vs-FFM contrast on both venues, we estimate the following regression:

$$\log y_{c,b} = \alpha_d + \gamma_b + \mu_h + \beta_{\text{NTM}} \mathbf{1}\{\text{NTM}_c\} + \delta_{\text{NTM}} \mathbf{1}\{b=29\} \cdot \mathbf{1}\{\text{NTM}_c\} + \varepsilon_{c,b},$$


where α_d , μ_h , γ_b capture the date, hour-of-day, and bin fixed effects. The standard errors are clustered at the cycle level. The FFM bin-29 increase is absorbed into the bin fixed effect γ_{29} , so δ_{NTM} measures the additional bin-29-vs-body activity increase in NTM cycles over the FFM baseline.

Table 4 reports the regression results. There are two key observations. First, on Binance, the bin-29 order flow is $\hat{\delta}_{\text{NTM}} = 1.37$ log points larger in NTM than in FFM cycles (about $e^{1.37} \approx 3.9$ times the FFM one on the typical cycle), and the same NTM concentration holds for absolute return (0.56), both positive at the 1% level. Polymarket diverges on flow: its NTM order-flow effect is *negative* (−0.18), the binary’s close spike being smaller in coin-flip cycles than in decided ones, where the one-sided sweep of the certain winner dominates. Its price-change effect is large and positive (2.62), but largely mechanical: resolution pulls a near-0.5 price to 0 or 1, while a decided contract’s price is already there. Where the outcome is in play, the incremental directional trade concentrates on the spot venue, not on the binary itself.

Second, Binance and Polymarket see opposite patterns in liquidity: into the NTM close Binance liquidity improves ($\hat{\delta}_{\text{NTM}} = 0.17$) while Polymarket liquidity worsens ($\hat{\delta}_{\text{NTM}} = -1.02$). The liquidity divergence is exactly the model’s prediction (Section 6.1). A maker quoting the Polymarket binary near the close becomes the counterparty to anyone about to push the spot across the strike: the manipulator buys the side he is about to make win, so the prediction-market maker who sold it to him is left on the losing side of the resolution. Polymarket makers therefore pull their book as the close nears, and pull it hardest in NTM cycles, where a flipping push is most likely. On Binance the



(a) Binance BTC spot market.



(b) Polymarket BTC 5-minute up/down contract.

Figure 10. Intra-cycle metrics per 5-min cycle, NTM vs. FFM (P3 only). This figure plots intra-cycle profiles on both venues, separately for NTM and FFM cycles. Each P3 5-minute cycle is classified NTM ($|p_{Up} - 0.5| < 0.10$) or FFM from the latest taker-focused on-chain Polymarket Up-token trade at time-to-maturity $TTM \in [10, 12]$ s; each metric is pooled over all classified P3 5-minute cycles. Each 1×3 grid plots one metric's intra-cycle profile against seconds into the 5-min cycle; the dotted vertical marks the final 10s before close. **(a) Binance BTC spot:** absolute order flow per cycle (\$M, |net signed taker dollar flow|), absolute 10s return ($|\Delta \log \text{mid}|$, bps), time-weighted depth within ± 10 ticks of the mid. **(b) Polymarket BTC 5-minute up/down contract:** absolute order flow per cycle (|net signed Up-token share flow|), absolute 10s probability change ($|\Delta p|$, pp), time-weighted depth within ± 5 cents of the mid (cumulative Up-token resting size at price levels within \$0.05 of the mid).

same push is uninformed and therefore generates only a *transitory* price impact. So the push flow is benign, and Binance makers supply *more* depth into the NTM close. Appendix C.1 refines the

Table 4. P3 near-settlement activity, NTM vs. FFM cycles, Binance and Polymarket. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and P3 cycles, $\log y_{c,b} = \alpha_d + \gamma_b + \mu_h + \beta_{\text{NTM}} \mathbf{1}\{\text{NTM}_c\} + \delta_{\text{NTM}} \mathbf{1}\{b=29\} \mathbf{1}\{\text{NTM}_c\} + \varepsilon_{c,b}$, with date, hour-of-day, and bin fixed effects. The first three columns report Binance spot, the last three the Polymarket 5-minute contract; within each venue the dependent variable $\log y_{c,b}$ is the log of the column metric. The activity column is *Order flow* on both venues — the per-cycle absolute net signed flow, in taker dollars on Binance ($|\sum_i d_i p_i a_i|$) and in Up-token shares on Polymarket. The second column per venue is *Abs. return* on Binance (absolute 10 s log return, $|\Delta \log \text{mid}|$, bps) and *Abs. price change* on Polymarket (absolute 10 s probability change, $|\Delta p|$, pp); the third is *Depth* (time-weighted resting size within a mid-anchored band: ± 10 ticks of the mid on Binance, ± 5 cents on Polymarket). A P3 cycle is NTM iff the 5m-contract Polymarket Up-token price is within 0.10 of 0.5 at $\text{TTM} \in [10, 12]$ s, FFM otherwise. δ_{NTM} is the bin-29-vs-body activity increase in NTM cycles over the FFM baseline (the latter absorbed in the bin FE); coefficients are log differences, so δ is a multiplicative effect of e^δ on the typical (geometric-mean) cycle. Appendix C.1 reports the same regression with the NTM cut refined into the full 15-min-overlap $\times$ per-contract-NTM split. One observation per (cycle, bin) cell. Polymarket order-book metrics (absolute return, depth) have smaller samples than order flow because the order-book archive begins later than the trade archive and outage days are excluded. HC1 standard errors in parentheses are clustered at the cycle level. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Binance			Polymarket		
	Order flow	Abs. return	Depth	Order flow	Abs. price change	Depth
δ_{NTM}	1.366*** (0.086)	0.555*** (0.052)	0.173*** (0.022)	−0.182*** (0.044)	2.618*** (0.089)	−1.016*** (0.052)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	395,790	307,108	395,790	395,154	138,405	142,468
Clusters (cycle)	15,223	15,221	15,223	15,223	5,631	5,631

NTM cut by 15-min contract overlap.

Post-settlement price reversal in NTM cycles. In addition to the activity spike, NTM cycles see a more pronounced post-settlement price reversal. Using the same moneyiness classifier, at each lag we estimate the following regression:

$$r_{j,c} = \alpha_d + \mu_h + \nu_g + \gamma_j r_{29,c} + \delta_{j,\text{NTM}} \mathbf{1}\{\text{NTM}_c\} r_{29,c} + \varepsilon_{j,c},$$

where γ_j is the FFM baseline reversal slope on r_{29} and $\delta_{j,\text{NTM}}$ is the NTM increment over it, so the NTM-cycle slope is $\gamma_j + \delta_{j,\text{NTM}}$ and $\delta_{j,\text{NTM}}$ is directly the NTM-minus-FFM reversal gap. α_d , μ_h , and ν_g are date, hour-of-day, and moneyiness-group fixed effects; standard errors are clustered by date.

We report the regression results in Table 5 and plot the slopes in Figure 11. FFM cycles revert by about 10% ($\gamma_{+1} = -0.10$), and NTM cycles revert significantly harder: the NTM increment is




Figure 11. Within-P3 post-settlement reversal slope γ_j , by 5m-contract moneyiness. This figure plots γ_j for each 5m-contract moneyiness group (the FFM baseline and the NTM total $\gamma_j + \delta_{j,NTM}$ of Table 5), on all P3 cycles split by 5m-contract Polymarket NTM (Up-token price within 0.10 of 0.5 near TTM = 10 s). Markers are γ_j ; error bars are $\pm 1.96 \times$ cluster-robust SE (clustered by date); $j = 0$ is omitted.

$\delta_{+1,NTM} = -0.15$, so NTM cycles revert by about 25% in total. A transient push pays off only by moving a price the prediction market has not already decided, so the cycles that revert hardest are exactly the NTM ones, where the push can flip the resolution. As before, in Appendix C.1 we further partition the cycles by 15-minute overlap. We find that the reversal is deeper in cycles where the 5m and 15m Polymarket contracts settle at the same close.

Table 5. Within-P3 post-settlement reversal slope by 5m-contract moneyiness. This table reports, at each lag $j \in \{-3, -2, -1, +1, +2, +3\}$, the two coefficients from a single cycle-level OLS on all P3 cycles (Feb 12 – Apr 8, 2026), $r_j = \alpha_d + \mu_h + \nu_g + \gamma_j r_{29} + \delta_{j,NTM} \mathbf{1}\{NTM\} r_{29} + \varepsilon$: γ_j is the FFM baseline reversal slope and $\delta_{j,NTM}$ the NTM increment over it, so the NTM-cycle slope is $\gamma_j + \delta_{j,NTM}$ and $\delta_{j,NTM}$ is the NTM-minus-FFM gap; the stars test each coefficient against zero. A cycle is NTM_{5m} iff the 5m-contract Polymarket Up-token price is within 0.10 of 0.5 near TTM $\in [10, 12]$ s, FFM otherwise. Date, hour-of-day, and moneyiness-group fixed effects; standard errors in parentheses are clustered by date. Negative j = same-cycle bin ($29 + j$); positive j = bin ($j - 1$) of the next cycle. Appendix C.1 dissects these groups further by 15-minute overlap. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
FFM _{5m}	-0.002 (0.011)	0.019 (0.014)	0.041*** (0.010)	-0.101*** (0.023)	-0.040** (0.018)	0.007 (0.013)
$\delta_{NTM_{5m}}$	0.008 (0.017)	-0.013 (0.024)	-0.004 (0.021)	-0.154*** (0.038)	0.028 (0.037)	-0.025 (0.026)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Cycle-group FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15,168					
Clusters (UTC date)	56					

6.3.3 The manipulation footprint fades at the fifteen-minute horizon

The model makes one more prediction: the manipulation footprint should weaken as the settlement horizon lengthens, because a longer window gives the asset's own price discovery more time to unfold, so a fixed-size push is less likely to be pivotal and manipulation opportunities become rarer *ex ante* (Section 6.1). The fifteen-minute contract provides the test. It went live earlier (in period P2), resolves based on the same Chainlink oracle price, and, during P2, was the most heavily traded of the crypto up/down contracts (Appendix B), so a weak footprint cannot be attributed to a small prize.

Identification. Within the five-minute clock, the fifteen-minute (and four-hour) contract settles only at the closes that coincide with a 15-minute mark: the *overlapping* cycles, one in three; the other closes carry no fifteen-minute settlement. The fifteen-minute effect is therefore identified by contrasting overlapping with non-overlapping closes. One complication forces a richer design: 15-minute marks might be intrinsically special on the spot market (round-time execution algorithms cluster there), so a raw overlap contrast would confound the fifteen-minute settlement with this clock effect. We difference it against the same contrast in the pre-Polymarket baseline: on the (cycle, bin) panel over P1 and P2, restricted to bins {0–24, 29} as before, we estimate

$$\log y_{c,b} = \theta_{t_c,b} + \alpha_d + \mu_h + \delta_{P2} \mathbf{1}\{b=29\} \cdot \mathbf{1}\{P2\} + \delta_{P2,5m \cap 15m} \mathbf{1}\{b=29\} \cdot \mathbf{1}\{P2\} \cdot D_{5m \cap 15m,c} + \varepsilon_{c,b},$$

where $D_{5m \cap 15m,c}$ indicates an overlapping cycle. The specification modifies the activity regression of Section 6.3.1 in two ways. First, the bin fixed effect is replaced by a cycle-type $\times$ bin fixed effect $\theta_{t_c,b}$, which gives overlapping and non-overlapping cycles each their own intra-cycle profile and thereby absorbs the clock effect, whatever its shape. Second, the bin-29 shift is split by overlap status: δ_{P2} is the bin-29-versus-body increase in P2 at non-overlapping closes, a pure period effect, and $\delta_{P2,5m \cap 15m}$ is the *additional* increase at overlapping closes, the difference-in-differences $(P2 - P1) \times (\text{overlap} - \text{non-overlap})$ that isolates the fifteen-minute settlement effect, netting out

both the period and the clock. The reversal regression of Section 6.3.1 is augmented in the same way. At each lag j , estimated separately on P1 and P2 cycles,

$$r_{j,c} = \alpha_d + \mu_h + \beta r_{29,c} + \beta_{5m \cap 15m} r_{29,c} D_{5m \cap 15m,c} \\ + \delta_{P2} r_{29,c} \mathbf{1}\{P2\} + \delta_{P2 \times (5m \cap 15m)} r_{29,c} \mathbf{1}\{P2\} \cdot D_{5m \cap 15m,c} + \varepsilon_{j,c},$$

with the lower-order level terms $D_{5m \cap 15m,c}$ and $\mathbf{1}\{P2\} \cdot D_{5m \cap 15m,c}$ included; the $\mathbf{1}\{P2\}$ level is absorbed by the date fixed effects. The four slopes mirror the activity design: β is the pre-Polymarket reversal slope at non-overlapping closes and $\beta_{5m \cap 15m}$ its overlap increment, the clock effect on the slope; δ_{P2} is the P2 slope change at non-overlapping closes, the period effect; and $\delta_{P2 \times (5m \cap 15m)}$ is the additional P2 change at overlapping closes, the difference-in-differences that gives the fifteen-minute settlement effect on the reversal slope.

Activity and reversal. Table 6 reports the activity difference-in-differences over P1 and P2, and Table 7 the reversal one. The fifteen-minute settlement increment $\delta_{P2, 5m \cap 15m}$ is small on every dimension: an order-flow increase of 0.068 against the five-minute 0.406, a depth increase of 0.021 against 0.048, no detectable absolute-return effect, and no post-settlement reversal at $j = +1$. Smaller but not exactly zero is what the horizon mechanism predicts: a longer window shrinks the incentive rather than eliminating it.

Conditioning on moneyness sharpens the picture but does not change it: even the fifteen-minute NTM cycles show at most a modest activity increase and no reversal, far below the five-minute footprint (Appendix C.2). The horizon prediction is borne out: the footprint that arrives in full force with the five-minute contract, the shortest, most frequent, and most cheaply moved settlement, appears at the fifteen-minute horizon only in strongly attenuated form.

Table 6. Fifteen-minute settlement: near-settlement activity, pre-Polymarket vs P2 by overlap. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and to the pre-Polymarket period (P1, Jul–Oct 2025) and the +15m/4h period (P2, Oct 2025 – Feb 2026). δ_{P2} is the bin-29-versus-body activity increase in P2 non-overlapping cycles relative to P1; $\delta_{P2, 5m \cap 15m}$ is the *additional* increase at overlapping closes (those landing on a 15-min mark, where the fifteen-minute contract settles) — the difference-in-differences $(P2 - P1) \times (\text{overlap} - \text{non-overlap})$ isolating the fifteen-minute settlement effect. Metrics and standard errors as in Table 2; fixed effects are date, hour-of-day, and cycle-type $\times$ bin. Coefficients are log differences. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Order flow	(2) Abs. return	(3) Depth
δ_{P2}	−0.062*** (0.017)	−0.091*** (0.019)	0.000 (0.003)
$\delta_{P2, 5m \cap 15m}$	0.068** (0.032)	0.042 (0.035)	0.021*** (0.005)
Date FE	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes
Cycle-type $\times$ Bin FE	Yes	Yes	Yes
Observations	1,684,789	1,071,666	1,684,782
Clusters (cycle)	64,800	64,391	64,800

Table 7. Fifteen-minute settlement: post-settlement reversal, pre-Polymarket vs P2 by overlap. This table reports, at each lag $j \in \{-3, \dots, +3\}$, a cycle-level OLS over P1 and P2 cycles. β is the pre-Polymarket reversal slope on r_{29} at non-overlapping closes, $\beta_{5m \cap 15m}$ its pre-Polymarket overlap increment (the clock effect at 15-min marks), δ_{P2} the P2 slope shift at non-overlapping closes, and $\delta_{P2 \times (5m \cap 15m)}$ the P2 overlap increment — the fifteen-minute settlement effect, a difference-in-differences. A negative slope is reversal. Date and hour-of-day fixed effects; HC1 standard errors in parentheses clustered by date. Negative j = same-cycle bin ($29 + j$); positive j = bin ($j - 1$) of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
β	−0.034 (0.047)	0.000 (0.018)	0.036 (0.022)	0.073*** (0.028)	0.000 (0.019)	0.010 (0.018)
$\beta_{5m \cap 15m}$	0.066 (0.047)	0.053* (0.030)	0.050 (0.034)	−0.081 (0.077)	−0.110** (0.045)	−0.039 (0.047)
δ_{P2}	0.033 (0.052)	−0.017 (0.033)	0.015 (0.030)	−0.116*** (0.036)	0.014 (0.024)	0.011 (0.025)
$\delta_{P2 \times (5m \cap 15m)}$	0.002 (0.085)	−0.058 (0.046)	−0.081* (0.043)	0.024 (0.087)	0.066 (0.060)	−0.005 (0.059)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	64,575					
Clusters (UTC date)	225					

6.4 Who wins and loses in manipulated cycles?

Section 6.3 shows that the five-minute launch coincides with a transitory final-10s spot push concentrated in the cycles the prediction market has left undecided: the manipulation footprint. This section tests the model's second prediction (Section 6.1): that the push transfers wealth, from the liquidity traders holding the losing side to the manipulator. We first single out the cycles carrying the largest manipulation pressure and document where they sit and whether the push succeeds; we then sort the traders active in them into types and measure who profits and who pays. Last, we rule out the alternative that the push is prediction-market makers hedging their binary exposure. Throughout this section, the sample is the five-minute-live period (P3), the same window as the within-period identification above.

6.4.1 Identification of manipulated cycles

We identify which cycles carry the largest manipulation pressure using a within-cycle, outcome-independent statistic:

$$PushIntensity_c = \frac{|\text{order flow}_{29,c}|}{\text{median}_{b \in \{0,\dots,24\}} |\text{order flow}_{b,c}|},$$

where $\text{order flow}_{b,c}$ is the Binance taker buy minus sell dollar volume in bin b of cycle c . The numerator is the magnitude of the final-10s order flow; the denominator is the typical magnitude of order flow per body bin in the *same* cycle. This normalization has three useful properties: (i) it is dimensionless and bounded below by zero; (ii) it is invariant to the overall scale of activity in the cycle, so a quiet and an active cycle with the same push-vs-body asymmetry score equally; and (iii) it is constructed entirely from the cycle's own intra-cycle Binance order flow, with no input from the resolution price, the Polymarket Up-token price, or the strike, so a trader's PnL is not mechanically related to its score.

We label a cycle a *manipulated cycle* if its *PushIntensity* is at or above the 90th percentile of the cycle distribution, $PushIntensity \geq 16.11$. Of the 16,073 cycles with on-chain Polymarket trading,

this gives 1,613 manipulated cycles and 14,460 *normal cycles*.¹⁹ The top-decile cut is a convention, chosen large enough that the manipulated-cycle sample carries meaningful statistical power (1,613 cycles $\times$ on average $\sim$ 1,500 traders per cycle) and small enough that the cycles selected are visibly distinct from the body of the distribution (the 90th-percentile *PushIntensity* of 16 is an order of magnitude larger than the median *PushIntensity* of $\sim$ 0.9, the latter close to what would be expected under no asymmetry between the final 10s and the body).

6.4.2 Characteristics of manipulated cycles

In Table 8, we contrast manipulated cycles with the normal ones along four dimensions: Binance activity, Polymarket activity, timing, and moneyness.

Because manipulated cycles are selected on *PushIntensity*, a large Binance spike among them is in part mechanical; the informative contrast is across venues. On Binance the bin-29 order flow averages \$1.7M against a typical body bin of just \$18k; on Polymarket the bin-29 order flow (\$2.8k) sits much closer to its own body bin (\$0.4k) and barely exceeds its normal-cycle level. The near-settlement push moves Binance spot without any comparable surge in Polymarket's own order flow, the first sign that the push is executed on the spot venue, matching the manipulation footprint.

Manipulated cycles are also *quieter* away from the close, not busier. The typical body-bin order flow is lower than in normal cycles on both venues, and the cycles fall disproportionately in thin-liquidity windows: 56% land in the Asia/overnight hours (vs 40% of normal cycles) and 44% on weekends (vs 27%), against only 23% in the deep European–US overlap (vs 35%). Yet despite this quieter backdrop the absolute Binance near-settlement order flow is an order of magnitude larger than in a normal cycle (\$1.7M vs \$68k). Manipulation thus concentrates precisely where the book is thinnest, the overnight and weekend hours, where a given dollar of near-settlement order flow moves the spot, and the Chainlink resolution it feeds, the most.

If the near-settlement push is manipulation, it should concentrate where the outcome is still in

¹⁹A further 55 cycles in the window, Polymarket outage periods and ultra-thin overnight cycles, carry a *PushIntensity* score from Binance order flow but have no Polymarket trades, and drop out of the trader-level analysis.

Table 8. Summary statistics for manipulated cycles vs normal cycles, P3 BTC 5m. This table reports summary statistics for manipulated cycles, the top decile by *PushIntensity* ($PushIntensity \geq 16.11$), against normal cycles, the remaining 90%. The *Binance activity* and *Polymarket activity* panels report the same two flow rows in thousands of dollars (USDC on Polymarket): the absolute net order flow in bin 29 (the final 10 s), and the per-cycle median of the absolute net order flow across the body bins (bins 0–24, the pre-ramp window, the same body denominator that enters *PushIntensity*). Bin 29 is a single 10 s bin and the body figure is the cycle’s typical body bin, so the two rows are on the same per-bin scale within each venue and across the two venues. *Binance activity* additionally reports the cross-cycle mean and median of *PushIntensity*. *Timing*: hour-of-day bucket (UTC) and weekday share. *Moneyness*: the share of cycles that are NTM (5m-contract Up-token price within 0.10 of 0.5 near TTM $\in [10, 12]$ s), computed over the cycles classifiable by that rule (99% of manipulated and 94% of normal cycles). All flow values are cross-cycle means.

	Manipulated cycles	Normal cycles
Number of cycles	1,613	14,460
<i>Binance activity</i>		
<i>PushIntensity</i> , mean	157.4	2.0
<i>PushIntensity</i> , median	57.8	0.9
Bin-29 signed notional (\$k)	\$1,714	\$68
Body signed notional , median/bin (\$k)	\$18	\$43
<i>Polymarket activity</i>		
Bin-29 signed notional (\$k)	\$2.83	\$2.69
Body signed notional , median/bin (\$k)	\$0.39	\$0.51
<i>Timing</i>		
Asia / overnight (21:00–07:00 UTC)	55.7%	39.9%
Europe, pre-US (07:00–13:00 UTC)	21.3%	25.5%
US & EU–US overlap (13:00–21:00 UTC)	23.1%	34.6%
Weekend (Sat–Sun)	43.6%	27.0%
<i>Moneyness</i>		
NTM share	18.7%	4.4%

play: the cycles in which moving the spot across the strike can actually flip the resolution. The *Moneyness* panel of Table 8 shows it does: among manipulated cycles 18.7% are NTM, against only 4.4% of normal cycles. Equivalently, a cycle that is NTM near the close is manipulated 33.0% of the time versus 9.1% for FFM cycles, a 3.6-fold higher rate. The two classifications are independent by construction: *PushIntensity* is built entirely from Binance signed flow and the moneyness cut entirely from the Polymarket Up-token price. So the concentration is not mechanical. It is the cross-sectional counterpart of the NTM activity and reversal results: the manipulation footprint is far more common exactly in the cycles where a push pays off.

6.4.3 Does manipulation succeed?

Does the push change the resolution? We classify each cycle by how confident Polymarket is just before the close: the favored-side price at TTM = 10 s, defined as $p_{\text{fav}} = \max(p_{\text{Up}}, 1 - p_{\text{Up}})$, ranges from 0.5 (a coin flip) to 1 (a settled outcome). Within manipulated cycles, the push can go in the *same* direction as the favored side or in the *opposite* direction, toward the underdog. Table 9 reports, by confidence bin and push direction, how often the side the push backs wins the resolution; the same side's win rate in normal cycles, which carry no push, is the baseline. Because manipulated cycles are selected on the size of the bin-29 flow, not on its direction or the outcome, the comparison is not mechanical.

The clearest pattern is that opposite-direction manipulation works: a push against the favored side raises the underdog's win rate far above its no-push baseline in every confidence bin, most dramatically where Polymarket is most confident. In the [0.90, 1.00) bin the underdog wins just 1.0% of normal cycles but 34.2% of manipulated ones: a push against a side the market priced at $\geq 90\%$ reverses the resolution roughly a third of the time, versus once in a hundred absent a push. Same-direction pushes, by contrast, matter only near the money: they add ~ 18 pp to the favored side's win rate in the near-NTM bin but are flat, even marginally negative, once the favored side already wins the large majority of normal cycles, where there is little left to push for.

The second pattern concerns the direction of the pushes: the split between same- and opposite-direction pushes is close to even in *every* confidence bin, including the most lopsided one, where Polymarket prices the favored side at $\geq 90\%$ yet half the pushes still drive against it. Manipulators are not merely reinforcing near-certain outcomes; they routinely bet on overturning them, and the bet is attractive: an underdog priced below ten cents that wins a third of the time returns several times its stake in expectation. We quantify the realized profits below.

Table 9. Manipulated-cycle outcome rates by Polymarket-favored confidence, P3 BTC 5m. This table reports, by confidence bin and push direction, how often the side the push backs wins the resolution. Cycles are binned by the Polymarket-favored confidence at $TTM = 10$ s, $p_{fav} = \max(p_{Up}, 1 - p_{Up})$ (width 0.10, folded over Up/Down since the empirical Up-vs-Down asymmetry is negligible). Each bin has two rows for the two manipulation directions, defined relative to the Polymarket-favored side: *Same-direction*, where the bin-29 Binance push agrees with the favored side, and *Opposite-direction*, where it pushes against the favored side (toward the underdog). *# cycles* and *% cycles* give the number and share of the bin’s manipulated cycles in each direction. *Win rate* is the rate at which the side the push targets, the favored side for same-direction, the underdog for opposite-direction, resolves as the winner among the manipulated cycles of that direction; *Win rate (normal cycles)* is the same side’s unconditional win rate in the bin’s normal cycles (which carry no push, hence no same/opposite split), serving as the no-manipulation baseline. Manipulated cycles are the top decile by *PushIntensity* ($PushIntensity \geq 16.11$); normal cycles are the remaining 90%.

p_{fav} range	Manipulation direction	# cycles	% cycles	Win rate	Win rate (normal cycles)
[0.50, 0.60)	Same-direction	149	52%	76.5%	58.6%
	Opposite-direction	136	48%	65.4%	41.4%
[0.60, 0.70)	Same-direction	125	46%	76.0%	70.9%
	Opposite-direction	145	54%	68.3%	29.1%
[0.70, 0.80)	Same-direction	111	54%	82.9%	82.3%
	Opposite-direction	94	46%	54.3%	17.7%
[0.80, 0.90)	Same-direction	95	43%	92.6%	91.3%
	Opposite-direction	125	57%	54.4%	8.7%
[0.90, 1.00)	Same-direction	300	50%	96.7%	99.0%
	Opposite-direction	304	50%	34.2%	1.0%

6.4.4 Trader types and PnL

PnL calculation. To attribute profits and losses across trader types, we compute realized PnL at the wallet-cycle level. Each on-chain Polymarket fill identifies the two counterparties by Polygon wallet address; for each (wallet, cycle) we accumulate the wallet’s signed cash flow and net Up/Down token positions from its fills, then settle the token positions at the Chainlink resolution price to obtain the cycle’s realized PnL:

$$PnL_{w,c} = \Delta cash_{w,c} + N_{Up,w,c} \mathbf{1}\{\text{Up wins}\} + N_{Down,w,c} \mathbf{1}\{\text{Down wins}\},$$

where $\Delta cash_{w,c}$ is the wallet’s USDC inflow during the cycle (positive when sells exceed buys, negative otherwise) and each token settles at \$1 for the winning side and \$0 for the losing side. Because each 5m market has its own token IDs, a wallet’s pre-cycle position is identically zero and each cycle is a closed loop. This is *gross* PnL: exchange fees are not subtracted; the fee

schedule changes mid-sample, and net-of-fee PnL is left to future work.²⁰ Settlement uses the Chainlink resolution price where the on-chain archive has it; for the small fraction of cycles with a missing Chainlink print, we fall back to the last trade in $[0, +30 \text{ s}]$ after close, where Polymarket secondary trade convergence to the Chainlink outcome takes $\sim 10 \text{ s}$. The panel covers 16,073 cycles with on-chain Polymarket trading and 243,155 distinct wallets, which we treat as 243,155 traders (multi-wallet identity by a single trader is unobserved).

Trader-type identification. We partition traders into three mutually exclusive types. A trader is labeled a *manipulator* if it satisfies two criteria across the sample: it participated in at least five manipulated cycles, and its aggregate net PnL across those cycles is at least \$2,000. The filter selects on participation and profitability in manipulated cycles, not on manipulation itself: a trader could in principle profit in those cycles by luck or by unrelated information.

A trader is labeled a *market maker* (MM) on the prediction-market contract if it does not pass the manipulator filter and satisfies three behavioral criteria, all defined on its trading across the sample. First, at least 85% of its fills have it on the resting (maker) side rather than crossing the spread. Second, conditional on being active in a cycle, it averages at least 50 fills per cycle, roughly one fill every six seconds. Third, its per-cycle median directional inventory at the close, defined as the absolute difference between its Up- and Down-token positions as a fraction of the gross shares it traded in that cycle, is at most 20%. The first two criteria capture passive execution and continuous quoting; the third separates liquidity provision from position-taking: an MM closes each cycle near delta-neutral rather than carrying a directional bet across resolution.

Every other trader is *retail/other*, the residual category. Under this classification the trader population breaks down to 821 manipulators (0.34% of the population), 227 MMs (0.09%), and 242,107 retail/other traders (99.57%).

²⁰Fee-enabled crypto contracts charge only the aggressor: resting liquidity pays no exchange fee and is eligible for rebates funded out of taker-fee revenue. The taker fee is proportional to $p(1 - p)$ in the contract price p , so it peaks at even odds, where short-horizon up/down contracts spend much of their lives. The schedule moved several times within our sample; the five-minute contracts were fee-enabled from launch, with a peak effective rate near 1.56% at even odds, raised to roughly 1.80% at the end of March 2026.

6.4.5 Profits and losses in manipulated cycles

Who profits and who pays in the manipulated cycles? Table 10 answers with four metrics per trader type, each split across manipulated and normal cycles: total net PnL, net PnL per cycle, the fraction of cycles in which the type loses money, and its share of the losses within a cycle. All four point to the same answer: the push transfers wealth from retail traders to manipulators, with market makers largely out of the way.

The PnL levels give the size of the transfer. Manipulators capture +\$8.2 M in aggregate across manipulated cycles, +\$5,096 per manipulated cycle. Retail traders fund nearly all of it: they lose −\$7.6 M, −\$4,715 per manipulated cycle. MMs largely dodge the transfer (−\$0.6 M, −\$381 per cycle), consistent with the Section 6.3.2 finding that MMs withdraw their book into the NTM close: they are rarely the counterparty because they do not quote in the cycles where the push is most likely.

The remaining two metrics show that the transfer is broad-based, not driven by a few large-loss cycles. Retail is net negative in 65% of manipulated cycles, against 48% of normal ones. The last metric works at the cycle level: the gross loss pool, defined as the sum across trader types of their net losses within a cycle, equals the sum of their gains by zero-sum, and a type's loss share is its loss divided by that pool. Weighted across cycles by pool size, retail bears 70.5% of manipulated-cycle losses versus 46.0% in normal cycles. The manipulator cohort still carries 20.2% of the manipulated-cycle pool even as the dominant net winner: that is the loss it books in the minority of cycles where its push fails, outweighed by its wins (in normal cycles, where it is more often a net loser, its share rises to 43.2%). The MM share is small in both regimes (9.3% and 10.8%): the loss flow runs from retail to manipulators directly.

Normal cycles show the reverse and confirm that the transfer is specific to manipulation. There, MMs collect the spread (+\$3.1 M, +\$215 per cycle), retail pays the ordinary cost of trading (−\$3.2 M, −\$221 per cycle), and manipulators break even (+\$0.1 M, +\$6 per cycle): the cohort that extracts \$5,096 per manipulated cycle has no edge in the remaining cycles. Retail's per-cycle loss is

Table 10. Trader-type PnL in manipulated vs normal cycles, P3 BTC 5m. This table reports four PnL metrics per trader type, split across manipulated and normal cycles. Trader types are mutually exclusive and exhaustive. *Manipulator*: traders with ≥ 5 manipulated cycles and $\geq \$2,000$ aggregate manipulated-cycle PnL. *MM*: residual of that filter satisfying passive-rate $\geq 85\%$, ≥ 50 fills per cycle when active, and per-cycle median directional inventory $\leq 20\%$ of gross shares traded. *Retail/other*: residual. Columns report four metrics, each split across manipulated (top-decile *PushIntensity*) and normal (other) cycles. *Total PnL* is the cohort's aggregate net PnL over all its (trader, cycle) observations of that type, in \$M; within each cycle type it sums to zero across cohorts by construction. *PnL per cycle* is Total PnL divided by the number of cycles of that type (1,613 manipulated, 14,460 normal). *% losing cycles* is the fraction of cycles in which the cohort is net negative, where the (cohort, cycle) cell aggregates a cohort's PnL across its member traders within a cycle. *Loss share* is the cohort's share of the *gross* loss pool (the sum across cohorts of their losses within a cycle), aggregated across cycles with pool-size weights, so it sums to 100% down each regime, and a net-winning cohort still carries a positive share from the cycles in which it happens to lose.

	Total PnL		PnL per cycle		% losing cycles		Loss share	
	Manip.	Normal	Manip.	Normal	Manip.	Normal	Manip.	Normal
Manipulator	+\$8.22M	+\$0.09M	+\$5,096	+\$6	38.0%	62.4%	20.2%	43.2%
Market maker	-\$0.62M	+\$3.11M	-\$381	+\$215	58.6%	37.7%	9.3%	10.8%
Retail / other	-\$7.61M	-\$3.20M	-\$4,715	-\$221	65.3%	47.7%	70.5%	46.0%

roughly twenty times larger when a push is present, and its counterparty switches from the MMs earning the spread to the manipulators. The full trader-type attribution is robust to the manipulator and MM thresholds (Appendix E).

6.4.6 Ruling out the gamma-hedging channel

The leading innocent explanation for the bin-29 push is hedging: a Polymarket maker who has sold Up tokens through the body offsets the resulting short by trading BTC, and that hedge, not manipulation, is what moves the spot across the strike. The channel requires no intent. Retail buys Up tokens through the body; makers absorb the demand by selling Up; the accumulated short must be hedged in the spot; the hedge pushes the close across the strike, and the body-period Up buyers profit. Two facts rule this out. The first, and the sharper, is the state of the contract when the push lands: in the cycles a push most often flips, the maker has essentially nothing to hedge. The second is the timing of the spot flow.

There is nothing to hedge in the cycles a push flips. A binary contract's sensitivity to the spot price is a delta that collapses away from the strike. When the favored side is priced near 0.5 the contract behaves almost like the spot itself, and a maker who is short it carries a real directional

exposure; but once the price moves toward 0 or 1, the contract is all but settled, its value barely responds to a further spot move, and the short position it represents is close to delta-flat. A maker holding such a position has almost nothing to hedge. Yet these near-certain cycles are exactly the ones the push overturns. In the $[0.90, 1.00)$ confidence bin — where Polymarket priced the favored side at $\geq 90\%$ ten seconds before the close, so the binary's delta has all but vanished — a push against the favored side reverses the resolution 34.2% of the time, against 1.0% in cycles with no push (Table 9). A trade placed when there is nothing to hedge cannot be a hedge. Whatever hedging response the spot flow contains can only come *after* the push, once it has dragged the price back toward the strike and revived the very exposure the hedging story needs; it cannot be the push's cause. The hedging channel is coherent only in the small near-the-money slice where the binary's delta is still live (about 6% of cycles); it has no purchase in the near-certain cycles that the push flips one time in three.

The timing of the spot flow is wrong for a hedge. The second fact concerns the cycles where there *is* something to hedge, and asks when the spot trading arrives. Under dynamic hedging, the standard practice in liquid markets, a maker hedges each marginal share as the short accumulates: if the cohort's short builds steadily through the body, Binance buying should ramp in step. Figure 12 shows it does not. The Polymarket positions build continuously — the manipulator cohort's net Up holdings climb to roughly 8,800 shares by the close, and the MM cohort absorbs a short that grows to roughly 1,350 shares as the final ten seconds begin, a near-linear accumulation rather than a flat book that dumps at the end. The matching spot hedge never appears through the body. Cumulative Binance order flow stands at just $-\$46\text{ k}$ at second 215 of the 300-second cycle, essentially zero and, if anything, slightly against the eventual push direction; the buying arrives only in the final minute, and the final ten seconds alone carry roughly $\$1.7\text{ M}$, essentially the cycle's entire net order flow in a single step. A position assembled only in the closing seconds is not a hedge accumulated against an exposure that built over the cycle; it is a single trade timed to move the settlement price.




Figure 12. Cumulative net position trajectory across the cycle, push-up vs. normal cycles, P3 BTC 5m. This figure plots, in each panel, the mean per-cycle cumulative position (or order flow) against seconds into the cycle, starting from zero at the open. The solid line conditions on *push-up* cycles, manipulated cycles whose near-settlement push is upward (labeled “manipulated cycles” in the legend), so that positions are not averaged across the two push directions; push-down cycles are reported in Appendix F. Panels (a)–(c) plot the net Up-token holdings (shares) of the manipulator, MM, and retail/other cohorts; panel (d) plots the cumulative Binance spot order flow (\$). Dashed gray lines: normal cycles as baselines. Pink shading: the final-10 s region (seconds 290–300).

The magnitudes rule out a deferred hedge. One variant survives the timing test: a maker who holds the naked delta through the body and hedges only as gamma surges at the very close would also trade in a single late burst. The magnitudes rule it out. A short binary position can lose at most its notional, so the MM cohort’s roughly 1,350-share short puts about \$1,400 per cycle at risk — three orders of magnitude below the \$1.7 M spot trade the channel would attribute to it. No maker, let alone all 28 active in the average manipulated cycle, hedges a four-figure maximum loss with a seven-figure spot position. The more credible reading is that the bin-29 push is direct spot-side action, coordinated with the manipulator cohort’s body-period positioning on Polymarket, with at most a residual hedging component in the bin-29 buying.

7 Conclusion

Prediction markets increasingly list contracts that settle on an asset price that holders can move by trading the underlying. A model and the on-chain record of Polymarket's five-minute Bitcoin contract agree: such a contract degrades price discovery in the underlying and transfers wealth from ordinary traders to the few who manipulate it, even as it makes the underlying more liquid.

The mechanism is a settlement-time push: a holder of the binary contract trades the underlying in the final seconds to move the settlement reference toward her side. The spot maker cannot separate this uninformed push from ordinary flow, so price discovery falls and wealth passes to the manipulator and, under a spread floor, to the prediction-market maker. The near-settlement spike appears only once the five-minute contract exists; the price move reverts within ten seconds; and the effect concentrates in the still-live cycles, where even a side priced above ninety percent is overturned about a third of the time, against one in a hundred otherwise. Some \$8.2 million was transferred over two months, mostly from retail.

The harm can be fixed: it tracks the contract horizon, the remedy our model provides. A manipulator enters only when a feasible push can carry the price across the strike cheaply enough, so a longer horizon makes pivotal price paths rarer and can eliminate the incentive. The settlement rule supplies parallel levers, from averaging or randomizing the close to capping the binary position or taxing settlement-window spot trades. Manipulation here is a settlement-design problem.

Appendix

A Omitted Proofs

A.1 Stage 2 Quotes and Posteriors

We prove Lemma 1 in this appendix.

Outcome probabilities. The informed trader enters with probability π . She buys if $V = +\sigma$ and sells if $V = -\sigma$ conditional on entry. A liquidity trader enters with probability $1 - \pi$. She buys or sells with equal probability conditional on entry. Therefore,

$$\mathbb{P}(Y_2^S = + \mid V = \pm\sigma) = \frac{1 \pm \pi}{2}, \quad \mathbb{P}(Y_2^S = - \mid V = \pm\sigma) = \frac{1 \mp \pi}{2}.$$

Marginalizing over V gives

$$q_+ = q_- = \frac{1}{2}.$$

Posterior and quotes. Bayes' rule gives

$$\rho_+ = \mathbb{P}(V = +\sigma \mid Y_2^S = +) = \frac{\mathbb{P}(Y_2^S = + \mid V = +\sigma) \mathbb{P}(V = +\sigma)}{q_+} = \frac{\frac{1+\pi}{2} \cdot \frac{1}{2}}{\frac{1}{2}} = \frac{1+\pi}{2},$$

$$\rho_- = 1 - \rho_+ = \frac{1-\pi}{2}.$$

Break-even pricing requires

$$\frac{1-\pi}{2} A_2^S = \frac{\pi}{2} (\sigma - A_2^S)$$

where the left-hand side is the expected profit from selling to a liquidity trader and the right-hand side is the expected loss from selling to the informed trader. Solving gives

$$A_2^S = \sigma\pi = \bar{V}_+.$$

By symmetry,

$$B_2^S = -\sigma\pi = \bar{V}_-.$$

The quotes are indeed inside $(-\sigma, \sigma)$ for $\pi \in (0, 1)$, confirming that the informed trader will never stay out at stage 2.

Posterior variance. Direct calculation of $\nu_{h_2} = 4\sigma^2\rho_{h_2}(1 - \rho_{h_2})$ gives

$$\nu_+ = \nu_- = \sigma^2(1 - \pi^2).$$

A.2 Stage 3 Quotes and Posteriors

We prove Proposition 1 in this appendix.

State-conditional outcome probabilities. The manipulator enters with probability p_M . She buys with probability α_{h_2} and stays out with probability $1 - \alpha_{h_2}$ conditional on entry. The informed trader enters with probability $(1 - p_M)\pi$. She buys if $V = +\sigma$ and sells if $V = -\sigma$ conditional on entry. A liquidity trader enters with probability $(1 - p_M)(1 - \pi)$. She buys or sells with equal probability conditional on entry. Therefore,

$$\begin{aligned}\mathbb{P}(Y_3^S = + \mid V = \pm\sigma, Y_2^S = h_2) &= p_M\alpha_{h_2} + (1 - p_M)\frac{1 \pm \pi}{2}, \\ \mathbb{P}(Y_3^S = 0 \mid V = \pm\sigma, Y_2^S = h_2) &= p_M(1 - \alpha_{h_2}), \\ \mathbb{P}(Y_3^S = - \mid V = \pm\sigma, Y_2^S = h_2) &= (1 - p_M)\frac{1 \mp \pi}{2}.\end{aligned}$$

Marginalizing over $V \mid Y_2^S = h_2$ gives

$$\begin{aligned} q_{h_2+} &= p_M \alpha_{h_2} + (1 - p_M) \left[\frac{1 - \pi}{2} + \pi \rho_{h_2} \right], \\ q_{h_20} &= p_M (1 - \alpha_{h_2}), \\ q_{h_2-} &= (1 - p_M) \left[\frac{1 - \pi}{2} + \pi (1 - \rho_{h_2}) \right]. \end{aligned}$$

Break-even ask. A stage-3 buy comes from the manipulator or a liquidity trader, both uninformed, or from the informed trader, who buys only when $V = +\sigma$. Break-even pricing requires

$$\left[p_M \alpha_{h_2} + (1 - p_M) \frac{1 - \pi}{2} \right] (A_{h_2} - \bar{V}_{h_2}) = (1 - p_M) \pi \rho_{h_2} (\sigma - A_{h_2})$$

where the left-hand side is the expected profit from selling to an uninformed buyer and the right-hand side is the expected loss from selling to the informed trader. Solving gives

$$A_{h_2} = \sigma \frac{p_M \alpha_{h_2} (2\rho_{h_2} - 1) + \frac{1-p_M}{2} [2\rho_{h_2} - 1 + \pi]}{p_M \alpha_{h_2} + \frac{1-p_M}{2} [1 + (2\rho_{h_2} - 1)\pi]}. \quad (4)$$

Substituting the stage-2 posteriors $\rho_+ = \frac{1+\pi}{2}$ and $\rho_- = \frac{1-\pi}{2}$ gives

$$A_+ = \frac{\sigma \pi [p_M \alpha_+ + (1 - p_M)]}{p_M \alpha_+ + (1 - p_M) \frac{1+\pi^2}{2}}, \quad A_- = \frac{-\sigma \pi p_M \alpha_-}{p_M \alpha_- + (1 - p_M) \frac{1-\pi^2}{2}}.$$

Break-even bid. The manipulator never sells, so a stage-3 sell comes only from a liquidity trader or from the informed trader, who sells only when $V = -\sigma$. Break-even pricing requires

$$(1 - p_M) \frac{1 - \pi}{2} (\bar{V}_{h_2} - B_{h_2}) = (1 - p_M) \pi (1 - \rho_{h_2}) (\sigma + B_{h_2})$$

where the left-hand side is the expected profit from buying from a liquidity seller and the right-hand side is the expected loss from buying from the informed trader. The common factor $(1 - p_M)$ cancels,

so the bid is independent of α_{h_2} and p_M . Solving gives

$$B_{h_2} = \sigma \frac{(2\rho_{h_2} - 1) - \pi}{1 - (2\rho_{h_2} - 1)\pi}. \quad (5)$$

Substituting the stage-2 posteriors $\rho_+ = \frac{1+\pi}{2}$ and $\rho_- = \frac{1-\pi}{2}$ gives

$$B_+ = 0, \quad B_- = \frac{-2\sigma\pi}{1 + \pi^2}.$$

No-trade posterior. No-trade occurs only when the manipulator is present, and she is uninformed. When it has positive probability ($\alpha_{h_2} < 1$) the posterior equals the prior entering stage 3; at $\alpha_{h_2} = 1$ the event has zero probability and is assigned the same belief by convention. Hence

$$\mathbb{P}(V = \pm\sigma \mid Y_3^S = \emptyset, Y_2^S = h_2) = \mathbb{P}(V = \pm\sigma \mid Y_2^S = h_2)$$

and

$$\bar{V}_{h_2\emptyset} = \bar{V}_{h_2} = \sigma(2\rho_{h_2} - 1),$$

constant in α_{h_2} and p_M .

Monotonicities. Recall the break-even pricing condition for the ask

$$\left[p_M \alpha_{h_2} + (1 - p_M) \frac{1 - \pi}{2} \right] (A_{h_2} - \bar{V}_{h_2}) = (1 - p_M) \pi \rho_{h_2} (\sigma - A_{h_2}).$$

For $p_M \in (0, 1)$ and $\pi \in (0, 1)$ both probability coefficients are strictly positive, so the balance forces $A_{h_2} \in (\bar{V}_{h_2}, \sigma)$, which we can also verify from the closed form above. The right-hand side does not contain α_{h_2} , while the left-hand side strictly increases in α_{h_2} . Raising α_{h_2} therefore moves the uninformed-profit side above the informed-loss side at the prevailing quote. Hence A_{h_2} has to strictly decrease to restore equality. For monotonicity in p_M , divide both sides by $1 - p_M > 0$ to get the

following, then apply a similar argument.

$$\left[\frac{p_M}{1 - p_M} \alpha_{h_2} + \frac{1 - \pi}{2} \right] (A_{h_2} - \bar{V}_{h_2}) = \pi \rho_{h_2} (\sigma - A_{h_2}).$$

The bid and the no-trade posterior mean are independent of both p_M and α_{h_2} , as the closed forms above show.

A.3 Selected Stage-3 Continuation Equilibrium

Definition 2 (Stable stage-3 fixed point). A stage-3 fixed point α_{h_2} is *stable* if for some $\eta > 0$,

$$(\alpha - \alpha_{h_2}) \Gamma_{h_2}(\alpha) \leq 0 \quad \text{for all } \alpha \in (\alpha_{h_2} - \eta, \alpha_{h_2} + \eta) \cap [0, 1]. \quad (6)$$

Equivalently, $\Gamma_{h_2} \geq 0$ just left and ≤ 0 just right of α_{h_2} . This excludes strict crossings from below, where local pressure pushes the conjecture away. Write $\mathcal{E}_{h_2}$ for the set of stable fixed points.

We prove Proposition 2 in this appendix. The closed-form expression for A_{h_2} shows that it is continuous in α_{h_2} , which implies that Γ_{h_2} is continuous in α_{h_2} as well since H is continuous. Moreover, $H(\cdot/\varepsilon)$ is constant outside $(-\varepsilon, \varepsilon)$ and affine inside, and $A_{h_2}(\alpha_{h_2})$ is continuous and strictly monotone, so Γ_{h_2} is strictly monotone on each of the at most three intervals on which the ask lies above, inside, or below $[-\varepsilon, \varepsilon]$.

Specifically, if the ask lies inside $[-\varepsilon, \varepsilon]$,

$$\Gamma'_{h_2}(\alpha_{h_2}) = \left(\frac{1}{2\varepsilon} - 1 \right) A'_{h_2}(\alpha_{h_2}) \neq 0$$

because $\varepsilon \neq \frac{1}{2}$.²¹ If the ask lies above ε or below $-\varepsilon$, then

$$\Gamma'_{h_2}(\alpha_{h_2}) = -A'_{h_2}(\alpha_{h_2}) > 0.$$

²¹The exclusion of the degenerate case $\varepsilon = \frac{1}{2}$ plays no role below, as the canonical regime of Proposition 4 has $\varepsilon < \sigma\pi < \frac{1}{2}$.

In particular, Γ_{h_2} has finitely many zeros. If $\Gamma_{h_2} \leq 0$ on a right-neighborhood of 0, then $\alpha_{h_2} = 0$ is stable. If $\Gamma_{h_2} \geq 0$ on a left-neighborhood of 1, then $\alpha_{h_2} = 1$ is stable. Otherwise neither holds, so every right-neighborhood of 0 contains a point with $\Gamma_{h_2} > 0$ and every left-neighborhood of 1 a point with $\Gamma_{h_2} < 0$; among the finitely many zeros lying between two such points, the first downcrossing has $\Gamma_{h_2} \geq 0$ immediately to its left and ≤ 0 immediately to its right, hence is stable. Therefore, there is at least one stable fixed point, which allows us to select the smallest $\alpha_{h_2}^* = \min \mathcal{E}_{h_2}$.

A.4 Price Discovery

We prove Proposition 3 in this appendix. By the law of total variance within each h_2 ,

$$\sum_{h_3} q_{h_2 h_3} \nu_{h_2 h_3} = \nu_{h_2} - \text{Var}(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2),$$

where

$$\text{Var}(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2) = \sum_{h_3} q_{h_2 h_3} (\bar{V}_{h_2 h_3} - \bar{V}_{h_2})^2$$

is the dispersion of the stage-3 posterior mean within each h_2 . Since stage 2 has no manipulator, its outcome probabilities q_{h_2} and conditional variances ν_{h_2} are all independent of p_M . Hence, it suffices to show, for each h_2 , that

$$\text{Var}(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2)(p_M) < \text{Var}(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2)(0).$$

By Proposition 1, $\bar{V}_{h_2 \emptyset}(p_M) = \bar{V}_{h_2}$, so the $h_3 = \emptyset$ term does not contribute to the within- h_2 dispersion. A sell order never comes from the manipulator, so the posterior mean after a sell is unaffected by p_M , but its probability is thinned by the manipulator's presence:

$$\bar{V}_{h_2-}(p_M) = \bar{V}_{h_2-}(0), \quad q_{h_2-}(p_M) = (1 - p_M) q_{h_2-}(0),$$

so the $h_3 = -$ term equals $(1 - p_M)$ times its benchmark value.

A buy order only moves the posterior mean when it comes from the informed trader:

$$\begin{aligned}
 & q_{h_2+}(p_M) \left(\bar{V}_{h_2+}(p_M) - \bar{V}_{h_2} \right) \\
 &= \mathbb{E} \left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{Y_3^S = +\} \mid Y_2^S = h_2 \right] \\
 &= \underbrace{\mathbb{E} \left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{Y_3^S = + \text{ from the manipulator} \} \mid Y_2^S = h_2 \right]}_{=0 \text{ (buy} \perp V)} \\
 &\quad + \underbrace{\mathbb{E} \left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{Y_3^S = + \text{ from a liquidity trader} \} \mid Y_2^S = h_2 \right]}_{=0 \text{ (buy} \perp V)} \\
 &\quad + \underbrace{\mathbb{E} \left[\left(V - \bar{V}_{h_2} \right) \mathbf{1}\{Y_3^S = + \text{ from the informed trader} \} \mid Y_2^S = h_2 \right]}_{=(1-p_M)\pi\rho_{h_2}(\sigma - \bar{V}_{h_2})} \\
 &= (1 - p_M) \pi \rho_{h_2} (\sigma - \bar{V}_{h_2}),
 \end{aligned}$$

which equals $(1 - p_M)$ times the benchmark quantity $q_{h_2+}(0) \left(\bar{V}_{h_2+}(0) - \bar{V}_{h_2} \right)$. Hence,

$$\begin{aligned}
 & q_{h_2+}(p_M) \left(\bar{V}_{h_2+}(p_M) - \bar{V}_{h_2} \right)^2 \\
 &= (1 - p_M)^2 \frac{q_{h_2+}(0)}{q_{h_2+}(p_M)} q_{h_2+}(0) \left(\bar{V}_{h_2+}(0) - \bar{V}_{h_2} \right)^2 \\
 &\leq (1 - p_M) q_{h_2+}(0) \left(\bar{V}_{h_2+}(0) - \bar{V}_{h_2} \right)^2
 \end{aligned}$$

because $q_{h_2+}(p_M) = p_M \alpha_{h_2} + (1 - p_M) q_{h_2+}(0) \geq (1 - p_M) q_{h_2+}(0)$.

Summing the three outcomes,

$$\text{Var} \left(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2 \right) (p_M) \leq (1 - p_M) \text{Var} \left(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2 \right) (0).$$

With $\pi > 0$ the informed trader moves the benchmark posterior, so

$$\text{Var} \left(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2 \right) (0) > 0.$$

Therefore, for any $p_M \in (0, 1)$,

$$\text{Var}(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2)(p_M) < \text{Var}(\bar{V}_{h_2 h_3} \mid Y_2^S = h_2)(0).$$

A.5 Canonical-Regime Equilibrium Calculations

We prove Proposition 4 in this appendix, which focuses on the canonical regime

$$\varepsilon < \sigma\pi < \frac{1}{2}.$$

A.5.1 Selected stage-3 continuation equilibrium

We find the selected continuation equilibrium at each price path by a three-step test of Proposition 2:

Step 1: Check whether $\alpha_{h_2} = 0$ is a stable fixed point.

Step 2: If $\alpha_{h_2} = 0$ fails, locate the interior fixed point by solving $\Gamma_{h_2}(\alpha_{h_2}) = 0$.

Step 3: If no interior fixed point exists either, the continuation equilibrium can only be $\alpha_{h_2} = 1$.

Favorable price path $h_2 = +$: pure stay-out equilibrium. For any $p_M \in (0, 1)$ and any $\alpha_+ \in [0, 1]$, $A_+(\alpha_+; p_M) \geq \bar{V}_+ = \sigma\pi > \varepsilon$, so $H(A_+(\alpha_+; p_M)/\varepsilon) = H(\bar{V}_+/\varepsilon) = 1$. Hence

$$\Gamma_+(\alpha_+; p_M) = 1 - 1 - [A_+(\alpha_+; p_M) - \sigma\pi] = -[A_+(\alpha_+; p_M) - \sigma\pi] \leq 0.$$

Therefore, for any $p_M \in (0, 1)$, $\alpha_+^* = 0$ is a stable fixed point and is clearly the smallest. Intuitively, after a favorable stage-2 buy the posterior already sits above the settlement band ($\bar{V}_+ = \sigma\pi > \varepsilon$), so the contract pays 1 whether or not the manipulator buys again; a stage-3 buy cannot raise her settlement payoff and only incurs spot cost, so she stays out.

Unfavorable price path $h_2 = -$: mixed or pure push. Fix $p_M \in (0, 1)$. At $\alpha_- = 0$, Proposition 1 gives $A_-(0; p_M) = 0$, while $\bar{V}_- = -\sigma\pi < -\varepsilon$. Therefore

$$\Gamma_-(0; p_M) = \frac{1}{2} - 0 - [0 - (-\sigma\pi)] = \frac{1}{2} - \sigma\pi > 0,$$

so $\alpha_- = 0$ is not a fixed point. For an interior fixed point with $A_- \in [-\varepsilon, \varepsilon]$, the fixed-point condition is

$$\frac{A_- + \varepsilon}{2\varepsilon} - (A_- + \sigma\pi) = 0,$$

which gives the candidate mixed-branch stage-3 ask

$$A_- = \frac{\varepsilon(2\sigma\pi - 1)}{1 - 2\varepsilon} \in (-\varepsilon, 0).$$

Inverting the closed-form expression in Proposition 1 gives the candidate mixed-branch intensity

$$\alpha_-(p_M) = \frac{1 - p_M}{p_M} \frac{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon)} > 0.$$

This candidate is admissible if and only if

$$\alpha_-(p_M) < 1 \iff p_M > p_0 := \frac{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon(1 - \pi^2)(1 - 2\sigma\pi)}.$$

Since A_- strictly decreases in α_- by Corollary 1, we can verify that Γ_- crosses zero from positive to negative as α_- crosses this candidate from below to above. Thus, for $p_M \in (p_0, 1)$, this candidate is the selected continuation equilibrium. For $p_M \in (0, p_0]$, this candidate would be equal to or greater than 1, which is not admissible, so the selected continuation equilibrium is the endpoint $\alpha_-^* = 1$:

$$\alpha_-^*(p_M) = \begin{cases} 1, & 0 < p_M \leq p_0, \\ \frac{1 - p_M}{p_M} \frac{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon)}, & p_0 < p_M < 1. \end{cases}$$

The corresponding equilibrium ask is

$$A_{-}^{*}(p_M) = \begin{cases} -\frac{2\sigma\pi p_M}{(1-\pi^2) + (1+\pi^2)p_M}, & 0 < p_M \leq p_0, \\ -\frac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon}, & p_0 < p_M < 1. \end{cases}$$

On both branches, this ask remains in the settlement band.²² Therefore the settlement probability induced by a stage-3 push after an unfavorable price path is

$$H(A_{-}^{*}(p_M)/\varepsilon) = \begin{cases} \frac{1}{2} - \frac{\sigma\pi p_M}{\varepsilon[(1-\pi^2) + (1+\pi^2)p_M]}, & 0 < p_M \leq p_0, \\ \frac{\sigma\pi - \varepsilon}{1-2\varepsilon}, & p_0 < p_M < 1. \end{cases}$$

The per-push spot cost after a stage-2 sell is

$$A_{-}^{*}(p_M) - \bar{V}_{-} = \begin{cases} \frac{\sigma\pi(1-\pi^2)(1-p_M)}{(1-\pi^2) + (1+\pi^2)p_M}, & 0 < p_M \leq p_0, \\ \frac{\sigma\pi - \varepsilon}{1-2\varepsilon}, & p_0 < p_M < 1. \end{cases}$$

The four displayed quantities are continuous at $p_M = p_0$, where the pure-push and mixed branches meet.

A.5.2 Optimal manipulator entry probability

Both the pricing of the binary contract and the manipulator's stage-1 entry problem use the selected continuation equilibrium from Appendix A.5.1. Combining Appendix A.1, the general ask and bid formulas (4) and (5), and Appendix A.5.1, the stage-3 asks, bids, and no-trade posterior means used

²²On the mixed branch, the ask lies in $(-\varepsilon, 0)$ by Assumption 1. On the pure-push branch, it is clearly below zero, and it is easy to verify that the ask is strictly decreasing in p_M . The two branches meet at $p_M = p_0$.

below are

$$A_+^* = \frac{2\sigma\pi}{1+\pi^2}, \quad A_-^*(p_M) = \begin{cases} -\frac{2\sigma\pi p_M}{(1-\pi^2) + (1+\pi^2)p_M}, & 0 < p_M \leq p_0, \\ -\frac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon}, & p_0 < p_M < 1, \end{cases}$$

$$B_+ = 0, \quad B_- = -\frac{2\sigma\pi}{1+\pi^2},$$

$$\bar{V}_+ = \sigma\pi, \quad \bar{V}_- = -\sigma\pi.$$

Pure-push branch $p_M \in (0, p_0]$. On this branch, $\alpha_+^* = 0$ and $\alpha_-^* = 1$. As a result, there is no spot cost after a stage-2 buy, while the spot cost after a stage-2 sell is $A_-^*(p_M) - \bar{V}_-$. The per-path expected binary-contract payoffs are:

$$\Pi_M^+(p_M, +) = \alpha_+^* H(A_+^*/\varepsilon) + (1 - \alpha_+^*) H(\bar{V}_+/\varepsilon) = H(\bar{V}_+/\varepsilon) = 1,$$

$$\Pi_M^+(p_M, -) = \alpha_-^* H(A_-^*(p_M)/\varepsilon) + (1 - \alpha_-^*) H(\bar{V}_-/\varepsilon) = H(A_-^*(p_M)/\varepsilon),$$

$$\Pi_L^+(p_M, +) = q_{++}(0) H(A_+^*/\varepsilon) + q_{+-}(0) H(B_+/\varepsilon) = \frac{1+\pi^2}{2} \cdot 1 + \frac{1-\pi^2}{2} \cdot \frac{1}{2} = \frac{3+\pi^2}{4},$$

$$\Pi_L^+(p_M, -) = q_{-+}(0) H(A_-^*(p_M)/\varepsilon) + q_{--}(0) H(B_-/\varepsilon) = \frac{1-\pi^2}{2} H(A_-^*(p_M)/\varepsilon),$$

where $q_{h_2 h_3}(0)$ denotes the conditional stage-3 outcome probabilities defined in Section 4.5 evaluated at $p_M = 0$.

The manipulator's expected binary-contract payoff net of expected spot manipulation cost is

$$\begin{aligned}
 & \Pi_M^+(p_M) - K^+(p_M) \\
 &= q_+ \Pi_M^+(p_M, +) + q_- \Pi_M^+(p_M, -) - q_+ \cdot 0 - q_- [A_-^*(p_M) - \bar{V}_-] \\
 &= \frac{1}{2} + \frac{1}{2} H(A_-^*(p_M)/\varepsilon) - \frac{1}{2} [A_-^*(p_M) - \bar{V}_-] \\
 &= \frac{1}{2} + \frac{1}{2} \cdot \frac{\varepsilon [(1 - \pi^2) + (1 + \pi^2) p_M] - 2\sigma\pi p_M}{2\varepsilon [(1 - \pi^2) + (1 + \pi^2) p_M]} - \frac{1}{2} \cdot \frac{\sigma\pi (1 - \pi^2) (1 - p_M)}{(1 - \pi^2) + (1 + \pi^2) p_M} \\
 &= \frac{1}{2} + \frac{(1 - \pi^2) (1 - 2\sigma\pi) (1 - p_M/p_0)}{4 [(1 - \pi^2) + (1 + \pi^2) p_M]}.
 \end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}
 & A_1^{P, BE}(p_M) \\
 &= p_M \Pi_M^+(p_M) + (1 - p_M) \Pi_L^+(p_M) \\
 &= p_M [q_+ \Pi_M^+(p_M, +) + q_- \Pi_M^+(p_M, -)] + (1 - p_M) [q_+ \Pi_L^+(p_M, +) + q_- \Pi_L^+(p_M, -)] \\
 &= p_M \left[\frac{1}{2} + \frac{1}{2} H(A_-^*(p_M)/\varepsilon) \right] + (1 - p_M) \left[\frac{3 + \pi^2}{8} + \frac{1 - \pi^2}{4} H(A_-^*(p_M)/\varepsilon) \right] \\
 &= \frac{p_M}{2} + (1 - p_M) \frac{3 + \pi^2}{8} + \left[\frac{p_M}{2} + \frac{(1 - p_M)(1 - \pi^2)}{4} \right] \left[\frac{1}{2} - \frac{\sigma\pi p_M}{\varepsilon [(1 - \pi^2) + (1 + \pi^2) p_M]} \right] \\
 &= \frac{1}{2} - \frac{p_M (\sigma\pi - \varepsilon)}{4\varepsilon} < \frac{1}{2}.
 \end{aligned}$$

Mixed branch $p_M \in (p_0, 1)$. On this branch, $\alpha_+^* = 0$ and $\alpha_-^*(p_M) \in (0, 1)$. The per-path expected binary-contract payoffs are:

$$\Pi_M^+(p_M, +) = 1 \quad (\text{same as the pure-push branch}),$$

$$\Pi_M^+(p_M, -) = \alpha_-^*(p_M) H(A_-^*(p_M)/\varepsilon) + (1 - \alpha_-^*(p_M)) H(\bar{V}_-/\varepsilon) = \alpha_-^*(p_M) H(A_-^*(p_M)/\varepsilon),$$

$$\Pi_L^+(p_M, +) = \frac{3 + \pi^2}{4} \quad (\text{same as the pure-push branch}),$$

$$\Pi_L^+(p_M, -) = q_{-+}(0) H(A_-^*(p_M)/\varepsilon) + q_{--}(0) H(B_-/\varepsilon) = \frac{1 - \pi^2}{2} H(A_-^*(p_M)/\varepsilon).$$

The manipulator's expected binary-contract payoff net of expected spot manipulation cost is

$$\begin{aligned} & \Pi_M^+(p_M) - K^+(p_M) \\ &= q_+ \Pi_M^+(p_M, +) + q_- \Pi_M^+(p_M, -) - q_+ \cdot 0 - q_- \alpha_-^*(p_M) [A_-^*(p_M) - \bar{V}_-] \\ &= \frac{1}{2} + \frac{\alpha_-^*(p_M) \sigma \pi - \varepsilon}{2} \frac{1}{1 - 2\varepsilon} - \frac{\alpha_-^*(p_M) \sigma \pi - \varepsilon}{2} \frac{1}{1 - 2\varepsilon} \\ &= \frac{1}{2}. \end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned} & A_1^{\text{P, BE}}(p_M) \\ &= p_M \Pi_M^+(p_M) + (1 - p_M) \Pi_L^+(p_M) \\ &= p_M [q_+ \Pi_M^+(p_M, +) + q_- \Pi_M^+(p_M, -)] + (1 - p_M) [q_+ \Pi_L^+(p_M, +) + q_- \Pi_L^+(p_M, -)] \\ &= p_M \left[\frac{1}{2} + \frac{\alpha_-^*(p_M) \sigma \pi - \varepsilon}{2} \frac{1}{1 - 2\varepsilon} \right] + (1 - p_M) \left[\frac{3 + \pi^2}{8} + \frac{(1 - \pi^2)(\sigma \pi - \varepsilon)}{4(1 - 2\varepsilon)} \right] \\ &= \frac{p_M}{2} + (1 - p_M) \frac{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)}{4(1 - 2\varepsilon)} + (1 - p_M) \left[\frac{3 + \pi^2}{8} + \frac{(1 - \pi^2)(\sigma \pi - \varepsilon)}{4(1 - 2\varepsilon)} \right] \\ &= \frac{1}{2} - \frac{(1 - p_M)(1 - \pi^2)(1 - 2\sigma\pi)}{8} < \frac{1}{2}. \end{aligned}$$

Non-crossing equilibrium ask and optimal entry probability. From the above calculations, $A_1^{P, BE}(p_M) < 1/2$ for any $p_M \in (0, 1)$. Therefore, the non-crossing equilibrium stage-1 ask is

$$A_1^{P*}(p_M) = \max \left\{ A_1^{P, BE}(p_M), \frac{1}{2} \right\} = \frac{1}{2}$$

for any $p_M \in (0, 1)$. The manipulator's expected profit conditional on entering at stage 1 is

$$S_M^+(p_M) = \Pi_M^+(p_M) - K^+(p_M) - \frac{1}{2},$$

which yields

$$S_M^+(p_M) = \begin{cases} \frac{(1 - \pi^2)(1 - 2\sigma\pi)(1 - p_M/p_0)}{4[(1 - \pi^2) + (1 + \pi^2)p_M]}, & 0 < p_M \leq p_0, \\ 0, & p_0 < p_M < 1. \end{cases}$$

The pure-push branch is positive when $0 < p_M < p_0$ and vanishes at $p_M = p_0$. Therefore, the manipulator maximizes the ex-ante profit

$$\bar{S}_M^+(p_M) = p_M S_M^+(p_M) = \frac{p_M (1 - \pi^2)(1 - 2\sigma\pi)(1 - p_M/p_0)}{4[(1 - \pi^2) + (1 + \pi^2)p_M]} \quad (7)$$

over $0 < p_M < p_0$. The first-order condition reduces to

$$(1 + \pi^2)p_M^2 + 2(1 - \pi^2)p_M - (1 - \pi^2)p_0 = 0.$$

The unique positive root is

$$p_M^* = \sqrt{\left(\frac{1 - \pi^2}{1 + \pi^2}\right)^2 + \frac{1 - \pi^2}{1 + \pi^2}p_0} - \frac{1 - \pi^2}{1 + \pi^2},$$

which we can verify is the unique maximizer over $(0, p_0)$. At p_M^* ,

$$\alpha_+^*(p_M^*) = 0,$$

$$\alpha_-^*(p_M^*) = 1,$$

$$S_M^+(p_M^*) = \frac{(1 - 2\sigma\pi) p_M^*}{4p_0},$$

$$\bar{S}_M^+(p_M^*) = \frac{(1 - 2\sigma\pi) p_M^{*2}}{4p_0}.$$

A.5.3 Comparative statics derivations

We verify the signs in Corollary 2 under the equilibrium of Proposition 4.

Pure-push cutoff p_0 . We can sign the partial derivatives by inspection of the closed-form expression for p_0

$$p_0 = \left(1 + \frac{2(\sigma\pi - \varepsilon)}{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)} \right)^{-1},$$

which gives

$$\frac{\partial p_0}{\partial \varepsilon} > 0, \quad \frac{\partial p_0}{\partial \sigma} < 0, \quad \frac{\partial p_0}{\partial \pi} < 0.$$

Optimal entry probability p_M^* . For $z \in (0, 1)$, $p_0 \in (0, 1)$, the map

$$(z, p_0) \mapsto \sqrt{z^2 + zp_0} - z = \frac{p_0}{\sqrt{1 + p_0/z} + 1}$$

strictly increases in both z and p_0 . Since

$$p_M^* = \sqrt{\left(\frac{1 - \pi^2}{1 + \pi^2} \right)^2 + \frac{1 - \pi^2}{1 + \pi^2} p_0} - \frac{1 - \pi^2}{1 + \pi^2},$$

and $(1 - \pi^2)/(1 + \pi^2)$ strictly decreases in π , we obtain

$$\frac{\partial p_M^*}{\partial \varepsilon} > 0, \quad \frac{\partial p_M^*}{\partial \sigma} < 0, \quad \frac{\partial p_M^*}{\partial \pi} < 0.$$

Stage-3 ask $A_-^*(p_M^*)$. Recall that

$$A_-^*(p_M^*) = -\frac{2\sigma\pi p_M^*}{(1-\pi^2) + (1+\pi^2)p_M^*} = -\frac{2\sigma\pi}{(1-\pi^2)/p_M^* + (1+\pi^2)} < 0.$$

This quantity depends on ε only through p_M^* and strictly decreases in p_M^* holding all other terms fixed. Since $\partial p_M^*/\partial \varepsilon > 0$, we obtain

$$\frac{\partial A_-^*(p_M^*)}{\partial \varepsilon} < 0.$$

For σ and π , using the closed-form expression for p_M^* , we have

$$(1-\pi^2) + (1+\pi^2)p_M^* = \sqrt{(1-\pi^2)^2 + (1-\pi^2)(1+\pi^2)p_0}$$

and therefore

$$A_-^*(p_M^*) = -\frac{2\sigma\pi}{1+\pi^2} \left[1 - \sqrt{\frac{1-\pi^2}{(1-\pi^2) + (1+\pi^2)p_0}} \right].$$

Substituting the closed-form expression for p_0 gives

$$\begin{aligned} A_-^*(p_M^*) &= -\frac{2\sigma\pi}{1+\pi^2} \left[1 - \sqrt{1 - \frac{\varepsilon(1+\pi^2)(1-2\sigma\pi)}{2\sigma\pi(1-2\varepsilon)}} \right] \\ &= -\frac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon} \left[1 + \sqrt{1 - \frac{\varepsilon(1+\pi^2)(1-2\sigma\pi)}{2\sigma\pi(1-2\varepsilon)}} \right]^{-1} \\ &= -\frac{\varepsilon(1-2\sigma\pi)}{1-2\varepsilon} \left[1 + \sqrt{1 - \frac{\varepsilon}{2\sigma(1-2\varepsilon)} \left[\pi + \frac{1}{\pi} - 2\sigma(1+\pi^2) \right]} \right]^{-1}. \end{aligned}$$

By inspection of the last expression, the square-root term strictly increases in both σ and π , while the positive fraction multiplying the leading negative sign strictly decreases in both σ and π . Therefore,

$$\frac{\partial A_-^*(p_M^*)}{\partial \sigma} > 0, \quad \frac{\partial A_-^*(p_M^*)}{\partial \pi} > 0.$$

Spot manipulation cost at p_M^* . Since $\bar{V}_- = -\sigma\pi$, the spot manipulation cost is

$$A_-^*(p_M^*) - \bar{V}_- = A_-^*(p_M^*) + \sigma\pi.$$

Therefore,

$$\begin{aligned}\frac{\partial}{\partial \varepsilon} [A_-^*(p_M^*) - \bar{V}_-] &= \frac{\partial A_-^*(p_M^*)}{\partial \varepsilon} < 0, \\ \frac{\partial}{\partial \sigma} [A_-^*(p_M^*) - \bar{V}_-] &= \frac{\partial A_-^*(p_M^*)}{\partial \sigma} + \pi > 0, \\ \frac{\partial}{\partial \pi} [A_-^*(p_M^*) - \bar{V}_-] &= \frac{\partial A_-^*(p_M^*)}{\partial \pi} + \sigma > 0.\end{aligned}$$

Ex-ante manipulation profit $\bar{S}_M^+(p_M^*)$. Using the first-order condition from Appendix A.5.2 evaluated at p_M^* , we can verify that

$$\bar{S}_M^+(p_M^*) = \frac{1 - 2\sigma\pi}{4} \frac{p_M^{*2}}{p_0} = \frac{1 - 2\sigma\pi}{4} \left(\frac{1 + \pi^2}{1 - \pi^2} + \frac{2}{p_M^*} \right)^{-1}.$$

Combined with the comparative statics results for p_M^* , we obtain by inspection of the last expression that

$$\frac{\partial \bar{S}_M^+(p_M^*)}{\partial \varepsilon} > 0, \quad \frac{\partial \bar{S}_M^+(p_M^*)}{\partial \sigma} < 0, \quad \frac{\partial \bar{S}_M^+(p_M^*)}{\partial \pi} < 0.$$

Conditional manipulation profit $S_M^+(p_M^*)$. Let

$$T := \sqrt{1 - \frac{\varepsilon(1 + \pi^2)(1 - 2\sigma\pi)}{2\sigma\pi(1 - 2\varepsilon)}} = \sqrt{\frac{1 - \pi^2}{(1 - \pi^2) + (1 + \pi^2)p_0}}.$$

In the canonical regime, $T \in (\sqrt{(1 - \pi^2)/2}, 1)$. Using the closed-form expression for p_M^* , we have

$$\frac{p_M^*}{p_0} = \frac{1 - \pi^2}{1 + \pi^2} \frac{1}{p_0} \left(\sqrt{1 + \frac{1 + \pi^2}{1 - \pi^2} p_0} - 1 \right) = \left(\sqrt{1 + \frac{1 + \pi^2}{1 - \pi^2} p_0} + 1 \right)^{-1} = \frac{T}{1 + T}.$$

Therefore, the conditional manipulation profit at the optimal entry probability decomposes as

$$S_M^+(p_M^*) = \underbrace{\frac{1 - 2\sigma\pi}{4}}_{\text{prize-compression channel}} \cdot \underbrace{\frac{T}{1 + T}}_{\text{entry-ratio channel}}.$$

Since $\partial p_0 / \partial \varepsilon > 0$, the expression for T shows that T strictly decreases in ε . Hence

$$\frac{\partial S_M^+(p_M^*)}{\partial \varepsilon} < 0.$$

For σ and π , the two channels move in opposite directions. The prize-compression channel $(1 - 2\sigma\pi) / 4$ strictly decreases in both σ and π . The entry-ratio channel $T / (1 + T)$ strictly increases in T , and direct differentiation of T gives

$$\begin{aligned} \frac{\partial T}{\partial \sigma} &= \frac{1 - T^2}{2\sigma T (1 - 2\sigma\pi)} > 0, \\ \frac{\partial T}{\partial \pi} &= \frac{(1 - T^2)(1 - \pi^2 + 4\sigma\pi^3)}{2\pi T (1 + \pi^2)(1 - 2\sigma\pi)} > 0. \end{aligned}$$

Thus the entry-ratio channel strictly increases in both σ and π . Putting the two channels together,

$$\begin{aligned} \frac{\partial S_M^+(p_M^*)}{\partial \sigma} &= \frac{1 - T - 4\sigma\pi T^2}{8\sigma T (1 + T)}, \\ \frac{\partial S_M^+(p_M^*)}{\partial \pi} &= \frac{(1 - T)(1 - \pi^2 + 4\sigma\pi^3) - 4\sigma\pi(1 + \pi^2)T^2}{8\pi T (1 + T)(1 + \pi^2)}. \end{aligned}$$

The denominators are strictly positive. Therefore,

$$\begin{aligned} \frac{\partial S_M^+(p_M^*)}{\partial \sigma} > 0 &\iff T < T_\sigma^* := \frac{2}{1 + \sqrt{1 + 16\sigma\pi}}, \\ \frac{\partial S_M^+(p_M^*)}{\partial \pi} > 0 &\iff T < T_\pi^* := \frac{2}{1 + \sqrt{1 + \frac{16\sigma\pi(1 + \pi^2)}{1 - \pi^2 + 4\sigma\pi^3}}}. \end{aligned}$$

Moreover, $T_\pi^* < T_\sigma^*$, because

$$1 + \pi^2 > 1 - \pi^2 + 4\sigma\pi^3 \iff \sigma\pi < \frac{1}{2}.$$

To state the sign conditions as settlement-noise cutoffs, note that $1 - T^2$ strictly decreases in T when $T \geq 0$, so for $j \in \{\sigma, \pi\}$,

$$\frac{\partial S_M^+(p_M^*)}{\partial j} > 0 \iff T < T_j^* \iff 1 - T^2 > 1 - T_j^{*2}.$$

Since

$$1 - T^2 = \frac{\varepsilon(1 + \pi^2)(1 - 2\sigma\pi)}{2\sigma\pi(1 - 2\varepsilon)}$$

strictly increases in ε , the last condition is equivalent to

$$\varepsilon > \varepsilon_j^* := \frac{2\sigma\pi(1 - T_j^{*2})}{(1 + \pi^2)(1 - 2\sigma\pi) + 4\sigma\pi(1 - T_j^{*2})},$$

provided $\varepsilon_j^* < \sigma\pi$. If $\varepsilon_j^* \geq \sigma\pi$, the partial derivative remains negative throughout the canonical regime. Since $T_\pi^* < T_\sigma^*$, we have $\varepsilon_\pi^* > \varepsilon_\sigma^*$. That is, a positive partial derivative with respect to π , when it exists, requires higher settlement noise than one with respect to σ .

When ε is small, both $\partial S_M^+/\partial\sigma$ and $\partial S_M^+/\partial\pi$ are negative: at low settlement noise the entry-ratio channel $T/(1 + T)$ is already near its ceiling of $1/2$, so increasing σ or π mainly compresses the prize and conditional profit falls.

A.5.4 Liquidity-trader loss decomposition

We prove Proposition 5. At $p_M^* \in (0, p_0)$, the selected continuation equilibrium after a stage-2 sell is on the pure-push branch. From Appendix A.5.2, the shadow stage-1 ask is

$$A_1^{\text{P,BE}}(p_M^*) = \frac{1}{2} - \frac{p_M^*(\sigma\pi - \varepsilon)}{4\varepsilon}.$$

Since the posted non-crossing ask is $A_1^{P*} = 1/2$, the prediction-market maker's expected rent per buy order is

$$R_P^{+*} := A_1^{P*} - A_1^{P, BE}(p_M^*) = \frac{p_M^* (\sigma\pi - \varepsilon)}{4\varepsilon} > 0.$$

The liquidity trader's payoff on the pure-push branch is

$$\Pi_L^+(p_M^*) = \frac{3 + \pi^2}{8} + \frac{1 - \pi^2}{4} H(A_-^*(p_M^*)/\varepsilon),$$

with

$$H(A_-^*(p_M^*)/\varepsilon) = \frac{1}{2} - \frac{\sigma\pi p_M^*}{\varepsilon[(1 - \pi^2) + (1 + \pi^2)p_M^*]}.$$

Therefore the expected loss per buy order for the liquidity trader is

$$L^+(p_M^*) := \frac{1}{2} - \Pi_L^+(p_M^*) = \frac{(1 - \pi^2)\sigma\pi p_M^*}{4\varepsilon[(1 - \pi^2) + (1 + \pi^2)p_M^*]} > 0.$$

Finally, the accounting identity follows directly from the definitions:

$$\begin{aligned} & R_P^{+*} + p_M^* S_M^+(p_M^*) + p_M^* K^+(p_M^*) \\ &= \frac{1}{2} - [p_M^* \Pi_M^+(p_M^*) + (1 - p_M^*) \Pi_L^+(p_M^*)] + p_M^* \left[\Pi_M^+(p_M^*) - K^+(p_M^*) - \frac{1}{2} \right] + p_M^* K^+(p_M^*) \\ &= (1 - p_M^*) \left[\frac{1}{2} - \Pi_L^+(p_M^*) \right] \\ &= (1 - p_M^*) L^+(p_M^*). \end{aligned}$$

A.6 Two-Ordinary-Round Extension Calculations

This appendix derives the two-ordinary-round extension used in Section 5.4. The analysis is conditional on a stage-1 prediction-market buy; the sell branch is symmetric.

A.6.1 Stage 2 Posteriors

Stage-2 price path probabilities. Stages 2.1 and 2.2 are independent ordinary Glosten–Milgrom spot rounds except that the prior over V entering stage 2.2 is the posterior from stage 2.1. At each round, the informed trader enters with probability π . She buys if $V = +\sigma$ and sells if $V = -\sigma$ conditional on entry. A liquidity trader enters with probability $1 - \pi$. She buys or sells with equal probability conditional on entry. Therefore, for $j \in \{2.1, 2.2\}$,

$$\mathbb{P}(Y_j^S = + \mid V = \pm\sigma) = \frac{1 \pm \pi}{2}, \quad \mathbb{P}(Y_j^S = - \mid V = \pm\sigma) = \frac{1 \mp \pi}{2}.$$

Conditional on V , the two ordinary spot orders are independent. Hence, for stage-2 price path $h_2 = h_{2,1}h_{2,2}$,

$$\mathbb{P}(h_2 = ++ \mid V = +\sigma) = \left(\frac{1 + \pi}{2}\right)^2, \quad \mathbb{P}(h_2 = ++ \mid V = -\sigma) = \left(\frac{1 - \pi}{2}\right)^2,$$

$$\mathbb{P}(h_2 = -- \mid V = +\sigma) = \left(\frac{1 - \pi}{2}\right)^2, \quad \mathbb{P}(h_2 = -- \mid V = -\sigma) = \left(\frac{1 + \pi}{2}\right)^2,$$

$$\mathbb{P}(h_2 = +- \mid V = +\sigma) = \mathbb{P}(h_2 = +- \mid V = -\sigma) = \frac{1 - \pi^2}{4},$$

$$\mathbb{P}(h_2 = -+ \mid V = +\sigma) = \mathbb{P}(h_2 = -+ \mid V = -\sigma) = \frac{1 - \pi^2}{4}.$$

Marginalizing over V gives

$$q_{++} = q_{--} = \frac{1 + \pi^2}{4}, \quad q_{+-} = q_{-+} = \frac{1 - \pi^2}{4}.$$

Posteriors. Bayes' rule gives

$$\rho_{++} = \mathbb{P}(V = +\sigma \mid h_2 = ++) = \frac{\mathbb{P}(h_2 = ++ \mid V = +\sigma) \mathbb{P}(V = +\sigma)}{q_{++}} = \frac{\left(\frac{1+\pi}{2}\right)^2 \cdot \frac{1}{2}}{\frac{1+\pi^2}{4}} = \frac{(1 + \pi)^2}{2(1 + \pi^2)}$$

and similarly

$$\rho_{+-} = \rho_{-+} = \frac{1}{2}, \quad \rho_{--} = \frac{(1 - \pi)^2}{2(1 + \pi^2)}.$$

Using $\bar{V}_{h_2} = \sigma(2\rho_{h_2} - 1)$ gives

$$\bar{V}_{++} = \frac{2\sigma\pi}{1 + \pi^2}, \quad \bar{V}_{+-} = \bar{V}_{-+} = 0, \quad \bar{V}_{--} = -\frac{2\sigma\pi}{1 + \pi^2}.$$

Posterior variance. Direct calculation of $v_{h_2} = 4\sigma^2\rho_{h_2}(1 - \rho_{h_2})$ gives

$$v_{++} = v_{--} = \frac{\sigma^2(1 - \pi^2)^2}{(1 + \pi^2)^2}, \quad v_{+-} = v_{-+} = \sigma^2.$$

A.6.2 Selected stage-3 continuation equilibrium

Recall that the gap between the expected continuation payoff from buying and that from staying out after stage-2 price path h_2 is

$$\Gamma_{h_2}(\alpha_{h_2}; p_M) = H\left(\frac{A_{h_2}(\alpha_{h_2}; p_M)}{\varepsilon}\right) - H\left(\frac{\bar{V}_{h_2}}{\varepsilon}\right) - [A_{h_2}(\alpha_{h_2}; p_M) - \bar{V}_{h_2}].$$

We find the selected continuation equilibrium at each stage-2 price path by the same three-step test of Appendix A.5.1. Throughout, stage-3 asks are obtained by substituting the relevant stage-2 posterior from Appendix A.6.1 into the general ask formula (4).

Favorable agreeing price path $h_2 = ++$: pure stay-out equilibrium. For any $p_M \in (0, 1)$ and any $\alpha_{++} \in [0, 1]$, the stage-3 ask is at least the prior mean, so

$$A_{++}(\alpha_{++}; p_M) \geq \bar{V}_{++} = \frac{2\sigma\pi}{1 + \pi^2} > \sigma\pi > \varepsilon.$$

Therefore $H(A_{++}(\alpha_{++}; p_M)/\varepsilon) = H(\bar{V}_{++}/\varepsilon) = 1$. Hence

$$\Gamma_{++}(\alpha_{++}; p_M) = 1 - 1 - [A_{++}(\alpha_{++}; p_M) - \bar{V}_{++}] \leq 0.$$

Thus, for any $p_M \in (0, 1)$, $\alpha_{++}^* = 0$ is a stable fixed point and is clearly the smallest.

Unfavorable agreeing price path $h_2 = --$: pure stay-out equilibrium. For any $p_M \in (0, 1)$,

$$\bar{V}_{--} = -\frac{2\sigma\pi}{1+\pi^2} < -\sigma\pi < -\varepsilon.$$

At $\alpha_{--} = 0$, substituting $\rho_{--} = (1 - \pi)^2 / (2(1 + \pi^2))$ into (4) gives

$$A_{--}(0; p_M) = -\sigma\pi < -\varepsilon.$$

Since the ask strictly decreases in α_{--} and is at least the prior mean,

$$\bar{V}_{--} \leq A_{--}(\alpha_{--}; p_M) \leq -\sigma\pi < -\varepsilon.$$

Consequently $H(A_{--}(\alpha_{--}; p_M) / \varepsilon) = H(\bar{V}_{--} / \varepsilon) = 0$, and

$$\Gamma_{--}(\alpha_{--}; p_M) = 0 - 0 - [A_{--}(\alpha_{--}; p_M) - \bar{V}_{--}] \leq 0.$$

Thus, for any $p_M \in (0, 1)$, $\alpha_{--}^* = 0$ is a stable fixed point and is clearly the smallest.

Disagreeing price paths $h_2 \in \{+-, -+\}$: pure push equilibrium. Fix $p_M \in (0, 1)$. For either disagreeing price path, $\bar{V}_{h_2} = 0$. At $\alpha_{h_2} = 0$, substituting $\rho_{h_2} = 1/2$ into (4) gives $A_{h_2}(0; p_M) = \sigma\pi$, and therefore

$$\Gamma_{h_2}(0; p_M) = 1 - \frac{1}{2} - [\sigma\pi - 0] = \frac{1}{2} - \sigma\pi > 0,$$

so $\alpha_{h_2} = 0$ is not a fixed point. For any $\alpha_{h_2} \in [0, 1]$, the ask is

$$A_{h_2}(\alpha_{h_2}; p_M) = \frac{\sigma\pi(1 - p_M)}{1 - p_M + 2p_M\alpha_{h_2}} \in \left[\frac{\sigma\pi(1 - p_M)}{1 + p_M}, \sigma\pi \right] \subset \left(0, \frac{1}{2} \right).$$

If $A_{h_2}(\alpha_{h_2}; p_M) \geq \varepsilon$, then

$$\Gamma_{h_2}(\alpha_{h_2}; p_M) = \frac{1}{2} - A_{h_2}(\alpha_{h_2}; p_M) \geq \frac{1}{2} - \sigma\pi > 0.$$

If $0 < A_{h_2}(\alpha_{h_2}; p_M) < \varepsilon$, then

$$\Gamma_{h_2}(\alpha_{h_2}; p_M) = A_{h_2}(\alpha_{h_2}; p_M) \left(\frac{1}{2\varepsilon} - 1 \right) > 0$$

because $\varepsilon < 1/2$. Thus no interior fixed point exists either, and the selected continuation equilibrium is the endpoint

$$\alpha_{+-}^* = \alpha_{-+}^* = 1.$$

A.6.3 Optimal manipulator entry probability

Both the pricing of the binary contract and the manipulator's stage-1 entry problem use the selected continuation equilibrium from Appendix A.6.2. Combining Appendix A.6.1, the general ask and bid formulas (4) and (5), and Appendix A.6.2, the stage-3 asks, bids, and no-trade posterior means used below are

$$A_{++}^*(p_M) = \frac{\sigma\pi(3 + \pi^2)}{1 + 3\pi^2}, \quad A_{--}^*(p_M) = -\sigma\pi, \quad A_{+-}^*(p_M) = A_{-+}^*(p_M) = \frac{\sigma\pi(1 - p_M)}{1 + p_M}.$$

$$B_{++} = \sigma\pi, \quad B_{--} = -\frac{\sigma\pi(3 + \pi^2)}{1 + 3\pi^2}, \quad B_{+-} = B_{-+} = -\sigma\pi.$$

$$\bar{V}_{++} = \frac{2\sigma\pi}{1 + \pi^2}, \quad \bar{V}_{+-} = \bar{V}_{-+} = 0, \quad \bar{V}_{--} = -\frac{2\sigma\pi}{1 + \pi^2}.$$

For each stage-2 price path h_2 , the per-path expected binary-contract payoffs are

$$\Pi_M^+(p_M, h_2) = \alpha_{h_2}^* H(A_{h_2}^*(p_M)/\varepsilon) + (1 - \alpha_{h_2}^*) H(\bar{V}_{h_2}/\varepsilon),$$

$$\Pi_L^+(p_M, h_2) = q_{h_2+}(0) H(A_{h_2}^*(p_M)/\varepsilon) + q_{h_2-}(0) H(B_{h_2}/\varepsilon).$$

Substituting the selected continuation equilibrium and the stage-3 quotes gives

$$\begin{aligned}\Pi_M^+(p_M, ++)&= 1, & \Pi_M^+(p_M, --)&= 0, \\ \Pi_M^+(p_M, +-)&= \Pi_M^+(p_M, -+)= H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right), \\ \Pi_L^+(p_M, ++)&= 1, & \Pi_L^+(p_M, --)&= 0, \\ \Pi_L^+(p_M, +-)&= \Pi_L^+(p_M, -+)= \frac{1}{2}H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right).\end{aligned}$$

Aggregating over the four stage-2 price paths gives

$$\begin{aligned}\Pi_M^+(p_M)&= \frac{1+\pi^2}{4} + \frac{1-\pi^2}{2}H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right), \\ \Pi_L^+(p_M)&= \frac{1+\pi^2}{4} + \frac{1-\pi^2}{4}H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right), \\ K^+(p_M)&= \frac{(1-\pi^2)\sigma\pi(1-p_M)}{2(1+p_M)}.\end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}A_1^{\text{P, BE}}(p_M) &= p_M \Pi_M^+(p_M) + (1-p_M) \Pi_L^+(p_M) \\ &= \frac{1+\pi^2}{4} + \frac{(1-\pi^2)(1+p_M)}{4} H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right) \\ &= \begin{cases} \frac{1}{2} + \frac{(1-\pi^2)p_M}{4}, & \frac{\sigma\pi(1-p_M)}{1+p_M} \geq \varepsilon, \\ \frac{1}{2} + \frac{(1-\pi^2)(1-p_M)(\sigma\pi-\varepsilon)}{8\varepsilon}, & \frac{\sigma\pi(1-p_M)}{1+p_M} < \varepsilon, \end{cases}\end{aligned}$$

which is strictly greater than 1/2 for any $p_M \in (0, 1)$. Therefore the non-crossing equilibrium stage-1

ask is just the break-even stage-1 ask:

$$A_1^{P*}(p_M) = \max \left\{ A_1^{P, BE}(p_M), \frac{1}{2} \right\} = A_1^{P, BE}(p_M)$$

for any $p_M \in (0, 1)$, and

$$B_1^{P*}(p_M) = 1 - A_1^{P*}(p_M).$$

The manipulator's expected profit conditional on entering at stage 1 is

$$\begin{aligned} S_M^+(p_M) &= \Pi_M^+(p_M) - K^+(p_M) - A_1^{P*}(p_M), \\ &= \frac{1 - \pi^2}{4} \left[(1 - p_M) H \left(\frac{\sigma\pi(1 - p_M)}{\varepsilon(1 + p_M)} \right) - \frac{2\sigma\pi(1 - p_M)}{1 + p_M} \right] \\ &= \frac{1 - \pi^2}{4} \begin{cases} \frac{(1 - p_M)(1 + p_M - 2\sigma\pi)}{1 + p_M}, & 0 < p_M \leq \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}, \\ \frac{(1 - p_M)[\varepsilon(1 + p_M) + \sigma\pi(1 - p_M) - 4\varepsilon\sigma\pi]}{2\varepsilon(1 + p_M)}, & \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon} < p_M < 1. \end{cases} \end{aligned}$$

Therefore, the manipulator maximizes the ex-ante profit

$$\bar{S}_M^+(p_M) = p_M S_M^+(p_M) = \frac{1 - \pi^2}{4} \Phi(p_M; \sigma, \pi, \varepsilon),$$

where

$$\Phi(p_M; \sigma, \pi, \varepsilon) := p_M \left[(1 - p_M) H \left(\frac{\sigma\pi(1 - p_M)}{\varepsilon(1 + p_M)} \right) - \frac{2\sigma\pi(1 - p_M)}{1 + p_M} \right]. \quad (8)$$

Equivalently,

$$\Phi(p_M; \sigma, \pi, \varepsilon) = \begin{cases} \frac{p_M(1 - p_M)(1 + p_M - 2\sigma\pi)}{1 + p_M}, & 0 < p_M \leq \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}, \\ \frac{p_M(1 - p_M)[\varepsilon(1 + p_M) + \sigma\pi(1 - p_M) - 4\varepsilon\sigma\pi]}{2\varepsilon(1 + p_M)}, & \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon} < p_M < 1. \end{cases}$$

The two branches agree at the cutoff $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, so Φ is continuous on $(0, 1)$. The first branch is strictly positive because $1 + p_M - 2\sigma\pi > 1 - 2\sigma\pi > 0$. On the second branch, the bracketed factor has slope $\varepsilon - \sigma\pi < 0$ in p_M , so it is bounded below by its value at $p_M = 1$, namely $2\varepsilon(1 - 2\sigma\pi) > 0$. Therefore, Φ is strictly positive on $(0, 1)$. If we define the boundary values

$$\Phi(0; \sigma, \pi, \varepsilon) = \lim_{p_M \downarrow 0^+} \Phi(p_M; \sigma, \pi, \varepsilon) = 0, \quad \Phi(1; \sigma, \pi, \varepsilon) = \lim_{p_M \uparrow 1^-} \Phi(p_M; \sigma, \pi, \varepsilon) = 0,$$

consistent with the convention that $\bar{S}_M^+(0) = \bar{S}_M^+(1) = 0$, then Φ is continuous on $[0, 1]$. Therefore, a maximizer exists and must be an interior point in $(0, 1)$. To show that it is unique, we can equivalently work with $\ln \Phi$, whose branchwise first-order derivative is

$$\frac{\partial \ln \Phi}{\partial p_M} = \begin{cases} \frac{1}{p_M} - \frac{1}{1 - p_M} + \frac{1}{1 + p_M - 2\sigma\pi} - \frac{1}{1 + p_M}, & 0 < p_M < \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}, \\ \frac{1}{p_M} - \frac{1}{1 - p_M} + \frac{\varepsilon - \sigma\pi}{\varepsilon(1 + p_M) + \sigma\pi(1 - p_M) - 4\varepsilon\sigma\pi} - \frac{1}{1 + p_M}, & \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon} < p_M < 1. \end{cases}$$

The branchwise second-order derivative is

$$\frac{\partial^2 \ln \Phi}{\partial p_M^2} = \begin{cases} -\frac{1}{p_M^2} - \frac{1}{(1 - p_M)^2} - \frac{1}{(1 + p_M - 2\sigma\pi)^2} + \frac{1}{(1 + p_M)^2} < 0, & 0 < p_M < \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}, \\ -\frac{1}{p_M^2} - \frac{1}{(1 - p_M)^2} - \frac{(\sigma\pi - \varepsilon)^2}{[\varepsilon + \sigma\pi - 4\varepsilon\sigma\pi - (\sigma\pi - \varepsilon)p_M]^2} + \frac{1}{(1 + p_M)^2} < 0, & \frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon} < p_M < 1. \end{cases}$$

On the first branch, the inequality follows because $0 < 1 + p_M - 2\sigma\pi < 1 + p_M$. On the second branch, the inequality follows because $0 < 1 - p_M < 1 + p_M$. Thus $\partial \ln \Phi / \partial p_M$ is strictly decreasing on both branches. At $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, the left derivative exceeds the right derivative:

$$\lim_{p_M \uparrow \left(\frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}\right)_-} \frac{\partial \ln \Phi}{\partial p_M} - \lim_{p_M \downarrow \left(\frac{\sigma\pi - \varepsilon}{\sigma\pi + \varepsilon}\right)_+} \frac{\partial \ln \Phi}{\partial p_M} = \frac{(\sigma\pi + \varepsilon)^2}{4\varepsilon\sigma\pi(1 - \sigma\pi - \varepsilon)} > 0.$$

Therefore, $\ln \Phi$ (and equivalently Φ) has a unique maximizer on $(0, 1)$. It is either the unique

interior zero of the first-order derivative or the kink $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$. Let $p_M^{*(2)}$ denote this unique maximizer. It remains to show that it is strictly greater than the corresponding maximizer in the baseline model, which we denote by $p_M^{*(1)}$.

Bounding the closed-form expression for $p_M^{*(1)}$ in Proposition 4, we have

$$0 < p_M^{*(1)} = \sqrt{\left(\frac{1 - \pi^2}{1 + \pi^2}\right)^2 + \frac{1 - \pi^2}{1 + \pi^2} p_0} - \frac{1 - \pi^2}{1 + \pi^2} < \sqrt{1 + \frac{\varepsilon(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon(1 - 2\sigma\pi)}} - 1 < \sqrt{2} - 1,$$

where we used the fact that $p_M^{*(1)}$ strictly increases in both $(1 - \pi^2) / (1 + \pi^2)$ and p_0 as established in Appendix A.5.3, $0 < (1 - \pi^2) / (1 + \pi^2) < 1$, and

$$0 < p_0 < \frac{\varepsilon(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon(1 - 2\sigma\pi)} < 1.$$

Because the first-order derivative of $\ln \Phi$ strictly decreases on each branch and jumps downward at $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, it suffices to show that it is still positive at the sharper upper bound on $p_M^{*(1)}$ displayed above. If that upper bound lies weakly below $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, only the first branch is relevant. It is strictly decreasing on $(0, 1)$ and equals $(\sqrt{2} - 2\sigma\pi)^{-1} > 0$ at $p_M = \sqrt{2} - 1$, so it must be positive at the sharper upper bound. If instead that upper bound lies strictly above $(\sigma\pi - \varepsilon) / (\sigma\pi + \varepsilon)$, it suffices to check the second branch at that value. At the sharper upper-bound value,

$$p_M^2 + 2p_M = \frac{\varepsilon(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon(1 - 2\sigma\pi)} \iff \varepsilon = \frac{2\sigma\pi(p_M^2 + 2p_M)}{(1 + 2\sigma\pi)(p_M^2 + 2p_M) + 1 - 2\sigma\pi}.$$

Substituting this expression into the third term in the second branch gives

$$\frac{\varepsilon - \sigma\pi}{\varepsilon(1 + p_M) + \sigma\pi(1 - p_M) - 4\varepsilon\sigma\pi} = \frac{p_M^2 + 2p_M - 1}{(1 + p_M)(p_M^2 + 4p_M + 1)}.$$

The remaining terms in the second branch satisfy

$$\frac{1}{p_M} - \frac{1}{1 - p_M} - \frac{1}{1 + p_M} = \frac{-p_M^2 - 2p_M + 1}{p_M(1 - p_M)(1 + p_M)}.$$

Therefore the second branch at the sharper upper-bound value equals

$$\frac{p_M^2 + 2p_M - 1}{1 + p_M} \left[\frac{1}{p_M^2 + 4p_M + 1} - \frac{1}{p_M(1 - p_M)} \right] = -\frac{(2p_M + 1)(p_M^2 + 2p_M - 1)}{p_M(1 - p_M)(p_M^2 + 4p_M + 1)} > 0,$$

with p_M evaluated at the sharper upper-bound value. The expression is positive because this sharper upper bound lies in $(0, \sqrt{2} - 1)$, so $p_M^2 + 2p_M - 1 < 0$ whereas all other factors are positive. Therefore, $\bar{S}_M^+(p_M)$ (under the two-ordinary-round extension) is still increasing at $p_M^{*(1)}$, which implies

$$p_M^{*(2)} > p_M^{*(1)}.$$

A.6.4 Price discovery

We prove the price-discovery part of Proposition 6. By the law of total variance, for $n = 1, 2$,

$$\mathbb{E} \left[v_{h_2 h_3}^{(n)} \right] (p_M) = v_\emptyset - \text{Var} \left(\bar{V}_{h_2 h_3}^{(n)} \right) (p_M) = \sigma^2 - \mathbb{E} \left[\left(\bar{V}_{h_2 h_3}^{(n)} (p_M) \right)^2 \right],$$

because $\mathbb{E} \left[\bar{V}_{h_2 h_3}^{(n)} (p_M) \right] = \mathbb{E} [V] = 0$ by the law of iterated expectations. As in Appendix A.4, lower residual variance is equivalent to greater dispersion of the final posterior mean.

We first compute the no-manipulation benchmarks. In the baseline model ($n = 1$), when $p_M = 0$, stages 2 and 3 consist of two ordinary spot rounds and thus share the same event structure as stage 2 under the two-ordinary-round extension, which allows us to reuse the results from Appendix A.6.1. The two agreeing price paths ($h_2 = h_3 = +$ or $-$) each have probability $(1 + \pi^2)/4$ and posterior means $\pm 2\sigma\pi/(1 + \pi^2)$, while the two disagreeing price paths ($h_2 = +, h_3 = -$ or $h_2 = -, h_3 = +$)

have final posterior mean zero. Hence

$$\mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (0) = \sigma^2 - 2 \cdot \frac{1 + \pi^2}{4} \left(\frac{2\sigma\pi}{1 + \pi^2} \right)^2 = \frac{\sigma^2 (1 - \pi^2)}{1 + \pi^2}.$$

Under the two-ordinary-round extension ($n = 2$), when $p_M = 0$, stages 2.1, 2.2, and 3 are three consecutive ordinary spot rounds. We can calculate the price path probabilities and posteriors following the same approach as in Appendix A.6 with an additional ordinary spot round. The all-buy and all-sell price paths each have probability $(1 + 3\pi^2)/8$ and posterior means $\pm\sigma\pi(3 + \pi^2)/(1 + 3\pi^2)$. The three ordered price paths with two buys and one sell have posterior mean $\sigma\pi$, the three ordered price paths with one buy and two sells have posterior mean $-\sigma\pi$, and each of these six mixed ordered price paths has probability $(1 - \pi^2)/8$. Therefore

$$\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0) = \sigma^2 - 2 \cdot \frac{1 + 3\pi^2}{8} \left(\frac{\sigma\pi(3 + \pi^2)}{1 + 3\pi^2} \right)^2 - 6 \cdot \frac{1 - \pi^2}{8} (\sigma\pi)^2 = \frac{\sigma^2 (1 - \pi^2)^2 (1 + 2\pi^2)}{1 + 3\pi^2}.$$

Hence the no-manipulation gain from adding the extra ordinary round is

$$\mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (0) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0) = \frac{\sigma^2 \pi^2 (1 - \pi^2) (1 + \pi^2 + 2\pi^4)}{(1 + \pi^2) (1 + 3\pi^2)} > 0.$$

We next compute the increase in $\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right]$ caused by the selected manipulation continuation for a fixed $p_M \in (0, 1)$. Applying the law of total variance to the dispersion of the final posterior mean as in Appendix A.4, we have

$$\text{Var} \left(\bar{v}_{h_2 h_3}^{(2)} \right) (p_M) = \text{Var} \left(\bar{v}_{h_2}^{(2)} \right) + \sum_{h_2} q_{h_2} \text{Var} \left(\bar{v}_{h_2 h_3}^{(2)} \mid Y_2^S = h_2 \right) (p_M).$$

$\text{Var} \left(\bar{v}_{h_2}^{(2)} \right)$ and q_{h_2} are independent of p_M . Therefore, the increase in $\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right]$ equals the weighted

loss of posterior-mean dispersion conditional on stage-2 price path h_2 :

$$\begin{aligned} & \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (p_M) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0) \\ &= \sum_{h_2} q_{h_2} \left[\text{Var} \left(\bar{v}_{h_2 h_3}^{(2)} \mid Y_2^S = h_2 \right) (0) - \text{Var} \left(\bar{v}_{h_2 h_3}^{(2)} \mid Y_2^S = h_2 \right) (p_M) \right]. \end{aligned}$$

At the two agreeing price paths $h_2 \in \{++, --\}$, $q_{++} = q_{--} = (1 + \pi^2)/4$, and, as we will prove below, the bracket equals

$$\frac{\sigma^2 \pi^2 p_M (1 - \pi^2)^3}{(1 + \pi^2)^2 (1 + 3\pi^2)}.$$

At the two disagreeing price paths $h_2 \in \{+-, -+\}$, $q_{+-} = q_{-+} = (1 - \pi^2)/4$, and, as we will prove below, the bracket equals

$$\frac{2\sigma^2 \pi^2 p_M}{1 + p_M}.$$

Therefore,

$$\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (p_M) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0) = \sigma^2 \pi^2 p_M (1 - \pi^2) \left[\frac{(1 - \pi^2)^2}{2(1 + \pi^2)(1 + 3\pi^2)} + \frac{1}{1 + p_M} \right],$$

which strictly increases in p_M by inspection.²³ Therefore, for every $p_M \in (0, 1)$,

$$\begin{aligned} & \mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (0) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (p_M) \\ &= \left[\mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (0) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0) \right] - \left[\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (p_M) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0) \right] \\ &= \sigma^2 \pi^2 (1 - \pi^2) (1 - p_M) \left[\frac{(1 - \pi^2)^2}{2(1 + \pi^2)(1 + 3\pi^2)} + \frac{1}{2(1 + p_M)} \right] \\ &> 0. \end{aligned}$$

²³It is straightforward to verify that the expression's left limit at $p_M = 1$ equals the gain from adding the extra ordinary round in the no-manipulation benchmark $\mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (0) - \mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (0)$, as expected, since when $p_M = 1$, stage 3 is fully occupied by the manipulator and does not contribute to price discovery at all.

Finally, since $p_M^{*(1)}, p_M^{*(2)} \in (0, 1)$,

$$\mathbb{E} \left[v_{h_2 h_3}^{(2)} \right] (p_M^{*(2)}) < \mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (0) < \mathbb{E} \left[v_{h_2 h_3}^{(1)} \right] (p_M^{*(1)}),$$

where the second inequality follows from Proposition 3. This proves the price-discovery part of Proposition 6.

We now verify the claims about $\text{Var}(\bar{V}_{h_2 h_3}^{(2)} | Y_2^S = h_2)(0) - \text{Var}(\bar{V}_{h_2 h_3}^{(2)} | Y_2^S = h_2)(p_M)$ below.

Agreeing price paths $h_2 \in \{++, --\}$. At $h_2 = ++$, the stage-2 posterior mean is $\bar{V}_{++} = 2\sigma\pi/(1 + \pi^2)$. When $p_M = 0$, stage 3 is another ordinary spot round, so the conditional probability of a stage-3 buy is

$$q_{+++}(0) = \rho_{++} \frac{1 + \pi}{2} + (1 - \rho_{++}) \frac{1 - \pi}{2} = \frac{1 + 3\pi^2}{2(1 + \pi^2)}$$

and the conditional probability of a stage-3 sell is

$$q_{++-}(0) = 1 - q_{+++}(0) = \frac{1 - \pi^2}{2(1 + \pi^2)}.$$

Bayes' rule gives the posterior mean after a stage-3 buy

$$\begin{aligned} & \bar{V}_{+++}^{(2)}(0) \\ &= \sigma \frac{\mathbb{P}(h_2 = ++, h_3 = + | V = +\sigma) \mathbb{P}(V = +\sigma) - \mathbb{P}(h_2 = ++, h_3 = + | V = -\sigma) \mathbb{P}(V = -\sigma)}{\mathbb{P}(h_2 = ++, h_3 = + | V = +\sigma) \mathbb{P}(V = +\sigma) + \mathbb{P}(h_2 = ++, h_3 = + | V = -\sigma) \mathbb{P}(V = -\sigma)} \\ &= \sigma \frac{\left(\frac{1+\pi}{2}\right)^3 \frac{1}{2} - \left(\frac{1-\pi}{2}\right)^3 \frac{1}{2}}{\left(\frac{1+\pi}{2}\right)^3 \frac{1}{2} + \left(\frac{1-\pi}{2}\right)^3 \frac{1}{2}} \\ &= \sigma \frac{(1 + \pi)^3 - (1 - \pi)^3}{(1 + \pi)^3 + (1 - \pi)^3} \\ &= \frac{\sigma\pi(3 + \pi^2)}{1 + 3\pi^2} \end{aligned}$$

and similarly the posterior mean after a stage-3 sell

$$\bar{V}_{++-}^{(2)}(0) = \sigma \frac{(1+\pi)^2(1-\pi) - (1-\pi)^2(1+\pi)}{(1+\pi)^2(1-\pi) + (1-\pi)^2(1+\pi)} = \sigma\pi.$$

Therefore the posterior-mean dispersion conditional on stage-2 price path $h_2 = ++$ is

$$\begin{aligned} & \text{Var}(\bar{V}_{++h_3}^{(2)} | Y_2^S = ++)(0) \\ &= \frac{1+3\pi^2}{2(1+\pi^2)} \left[\frac{\sigma\pi(3+\pi^2)}{1+3\pi^2} - \frac{2\sigma\pi}{1+\pi^2} \right]^2 + \frac{1-\pi^2}{2(1+\pi^2)} \left[\sigma\pi - \frac{2\sigma\pi}{1+\pi^2} \right]^2 \\ &= \frac{\sigma^2\pi^2(1-\pi^2)^3}{(1+\pi^2)^2(1+3\pi^2)}. \end{aligned}$$

When $p_M \in (0, 1)$, the selected continuation equilibrium has $\alpha_{++}^* = 0$, so the manipulator does not pool with ordinary buy or sell orders. When she occupies stage 3 with probability p_M , she always produces a no-trade outcome which does not move the posterior and hence contributes zero to $\text{Var}(\bar{V}_{++h_3}^{(2)} | Y_2^S = ++)(p_M)$. Otherwise, with probability $1 - p_M$, the manipulator is absent from stage 3 and it becomes an ordinary spot round. Thus

$$\text{Var}(\bar{V}_{++h_3}^{(2)} | Y_2^S = ++)(p_M) = (1 - p_M) \frac{\sigma^2\pi^2(1-\pi^2)^3}{(1+\pi^2)^2(1+3\pi^2)}.$$

The same calculation applies at $h_2 = --$ by symmetry:

$$\text{Var}(\bar{V}_{--h_3}^{(2)} | Y_2^S = --)(p_M) = (1 - p_M) \frac{\sigma^2\pi^2(1-\pi^2)^3}{(1+\pi^2)^2(1+3\pi^2)}.$$

Subtracting the two expressions verifies the claim.

Disagreeing price paths $h_2 \in \{+-, -+\}$. At $h_2 = +-$, the stage-2 posterior mean is $\bar{V}_{+-} = 0$.

When $p_M = 0$, stage 3 is another ordinary spot round, so the conditional probability of a stage-3 buy

is

$$q_{+-+}(0) = \rho_{+-} \frac{1+\pi}{2} + (1-\rho_{+-}) \frac{1-\pi}{2} = \frac{1}{2}$$

and the conditional probability of a stage-3 sell is

$$q_{+--}(0) = 1 - q_{+-+}(0) = \frac{1}{2}.$$

Bayes' rule gives the posterior mean after a stage-3 buy

$$\begin{aligned} & \bar{V}_{+-+}^{(2)}(0) \\ &= \sigma \frac{\mathbb{P}(h_2 = +-, h_3 = + | V = +\sigma) \mathbb{P}(V = +\sigma) - \mathbb{P}(h_2 = +-, h_3 = + | V = -\sigma) \mathbb{P}(V = -\sigma)}{\mathbb{P}(h_2 = +-, h_3 = + | V = +\sigma) \mathbb{P}(V = +\sigma) + \mathbb{P}(h_2 = +-, h_3 = + | V = -\sigma) \mathbb{P}(V = -\sigma)} \\ &= \sigma \frac{\left(\frac{1+\pi}{2}\right)^2 \frac{1-\pi}{2} \frac{1}{2} - \left(\frac{1-\pi}{2}\right)^2 \frac{1+\pi}{2} \frac{1}{2}}{\left(\frac{1+\pi}{2}\right)^2 \frac{1-\pi}{2} \frac{1}{2} + \left(\frac{1-\pi}{2}\right)^2 \frac{1+\pi}{2} \frac{1}{2}} \\ &= \sigma \frac{(1+\pi)^2 (1-\pi) - (1-\pi)^2 (1+\pi)}{(1+\pi)^2 (1-\pi) + (1-\pi)^2 (1+\pi)} \\ &= \sigma\pi, \end{aligned}$$

and similarly the posterior mean after a stage-3 sell

$$\bar{V}_{+--}^{(2)}(0) = \sigma \frac{(1+\pi)(1-\pi)^2 - (1-\pi)(1+\pi)^2}{(1+\pi)(1-\pi)^2 + (1-\pi)(1+\pi)^2} = -\sigma\pi.$$

Therefore the posterior-mean dispersion conditional on stage-2 price path $h_2 = +-$ is

$$\text{Var}\left(\bar{V}_{+-h_3}^{(2)} \mid Y_2^S = +- \right)(0) = \frac{1}{2} [\sigma\pi - 0]^2 + \frac{1}{2} [-\sigma\pi - 0]^2 = \sigma^2 \pi^2.$$

When $p_M \in (0, 1)$, the selected continuation equilibrium has $\alpha_{+-}^* = 1$ and equilibrium ask and bid as given in Appendix A.6.3:

$$A_{+-}^*(p_M) = \frac{\sigma\pi(1 - p_M)}{1 + p_M}, \quad B_{+-} = -\sigma\pi.$$

The conditional probability of a stage-3 buy is

$$q_{++}(p_M) = p_M + (1 - p_M) q_{++}(0) = \frac{1 + p_M}{2}$$

and the conditional probability of a stage-3 sell is

$$q_{--}(p_M) = 1 - q_{++}(p_M) = \frac{1 - p_M}{2}.$$

A no-trade outcome at stage-3 must come from the manipulator, so it does not move the posterior and contributes zero to $\text{Var}(\tilde{V}_{+-h_3}^{(2)} | Y_2^S = +-)(p_M)$. Therefore,

$$\text{Var}(\tilde{V}_{+-h_3}^{(2)} | Y_2^S = +-)(p_M) = \frac{1 + p_M}{2} \left[\frac{\sigma\pi(1 - p_M)}{1 + p_M} - 0 \right]^2 + \frac{1 - p_M}{2} [-\sigma\pi - 0]^2 = \sigma^2\pi^2 \frac{1 - p_M}{1 + p_M}.$$

The same calculation applies at $h_2 = -+$ by symmetry:

$$\text{Var}(\tilde{V}_{-+h_3}^{(2)} | Y_2^S = -+)(p_M) = \sigma^2\pi^2 \frac{1 - p_M}{1 + p_M}.$$

Subtracting the two expressions verifies the claim.

A.7 *n*-Ordinary-Round Extension Calculations

We prove Proposition 7 in this appendix. The analysis is conditional on a stage-1 prediction-market buy; the sell branch is symmetric.

A.7.1 Stage 2 Posteriors

Fix $n \geq 1$. A stage-2 price path is a sequence of n outcomes:

$$h_2 = h_{2,1} \cdots h_{2,n} \in \{+, -\}^n.$$

As we will see, the relevant sufficient statistic for the stage-2 posterior that enters stage 3 is the order imbalance

$$d(h_2) = \sum_{j=1}^n \mathbf{1}\{h_{2,j} = +\} - \sum_{j=1}^n \mathbf{1}\{h_{2,j} = -\} = 2 \sum_{j=1}^n \mathbf{1}\{h_{2,j} = +\} - n,$$

which takes values in the set

$$\mathcal{D}_n = \{-n, -n+2, \dots, n-2, n\}.$$

For $k \in \mathcal{D}_n$, let

$$b(k) = \frac{n+k}{2}, \quad s(k) = \frac{n-k}{2},$$

so that stage-2 price paths with imbalance k have $b(k)$ buys and $s(k)$ sells. Conditional on $V = +\sigma$, any individual stage-2 price path with imbalance k has probability

$$\left(\frac{1+\pi}{2}\right)^{b(k)} \left(\frac{1-\pi}{2}\right)^{s(k)},$$

whereas conditional on $V = -\sigma$, it has probability

$$\left(\frac{1-\pi}{2}\right)^{b(k)} \left(\frac{1+\pi}{2}\right)^{s(k)}.$$

Averaging over the symmetric prior on V , any individual stage-2 price path h_2 with $d(h_2) = k$ has probability

$$q_{h_2} = \frac{(1+\pi)^{b(k)} (1-\pi)^{s(k)} + (1-\pi)^{b(k)} (1+\pi)^{s(k)}}{2^{n+1}}.$$

Thus the total probability of all stage-2 price paths with imbalance k is

$$q_{d=k} := \sum_{h_2: d(h_2)=k} q_{h_2} = \binom{n}{b(k)} \frac{(1+\pi)^{b(k)} (1-\pi)^{s(k)} + (1-\pi)^{b(k)} (1+\pi)^{s(k)}}{2^{n+1}}. \quad (9)$$

Since $b(-k) = s(k) = n - b(k)$ and similarly $s(-k) = b(k) = n - s(k)$, we can verify that

$$q_{d=-k} = q_{d=k}. \quad (10)$$

Let

$$r := \frac{1+\pi}{1-\pi}$$

denote the likelihood ratio of a buy order in an ordinary spot round: a buy is r times as likely when $V = +\sigma$ as when $V = -\sigma$. By symmetry, the likelihood ratio of a sell order is r^{-1} . Since the n ordinary spot rounds are conditionally independent given V , Bayes' rule gives directly, for any stage-2 price path h_2 with $d(h_2) = k$,

$$\begin{aligned} \rho_{h_2} &= \mathbb{P}(V = +\sigma \mid Y_2^S = h_2) \\ &= \frac{\mathbb{P}(Y_2^S = h_2 \mid V = +\sigma) \mathbb{P}(V = +\sigma)}{\mathbb{P}(Y_2^S = h_2 \mid V = +\sigma) \mathbb{P}(V = +\sigma) + \mathbb{P}(Y_2^S = h_2 \mid V = -\sigma) \mathbb{P}(V = -\sigma)} \\ &= \frac{\left(\frac{1+\pi}{2}\right)^{b(k)} \left(\frac{1-\pi}{2}\right)^{s(k)} \frac{1}{2}}{\left(\frac{1+\pi}{2}\right)^{b(k)} \left(\frac{1-\pi}{2}\right)^{s(k)} \frac{1}{2} + \left(\frac{1-\pi}{2}\right)^{b(k)} \left(\frac{1+\pi}{2}\right)^{s(k)} \frac{1}{2}} = \frac{r^k}{1 + r^k}, \end{aligned}$$

where the last equality cancels common factors and uses $b(k) - s(k) = k$. This shows that ρ_{h_2} depends on h_2 only through $k = d(h_2)$, which proves that $d(h_2)$ is indeed a sufficient statistic for the stage-2 posterior. We write this common value as

$$\rho_{d=k} = \frac{r^k}{1 + r^k}.$$

The posterior mean and variance are therefore

$$\bar{V}_{d=k} = \sigma (2\rho_{d=k} - 1) = \sigma \frac{r^k - 1}{r^k + 1},$$

and

$$\nu_{d=k} = 4\sigma^2 \rho_{d=k} (1 - \rho_{d=k}) = \frac{4\sigma^2 r^k}{(1 + r^k)^2}.$$

Thus every stage-2 price path with the same imbalance induces the same posterior:

$$\rho_{h_2} = \rho_{d=d(h_2)}, \quad \bar{V}_{h_2} = \bar{V}_{d=d(h_2)}, \quad \nu_{h_2} = \nu_{d=d(h_2)}.$$

A.7.2 Selected stage-3 continuation equilibrium

Substituting $\rho_{d=k}$ into the general ask and bid formulas (4) and (5) gives

$$A_{d=k}(\alpha_{d=k}; p_M) = \sigma \frac{p_M \alpha_{d=k} \left(\frac{r^k - 1}{r^k + 1} \right) + \frac{1 - p_M}{2} \left(\frac{r^k - 1}{r^k + 1} + \pi \right)}{p_M \alpha_{d=k} + \frac{1 - p_M}{2} \left(1 + \frac{r^k - 1}{r^k + 1} \pi \right)}, \quad B_{d=k} = \sigma \frac{r^{k-1} - 1}{r^{k-1} + 1}.$$

At imbalance k , the gap between the expected continuation payoff from buying and that from staying out is

$$\Gamma_{d=k}(\alpha_{d=k}; p_M) = H\left(\frac{A_{d=k}(\alpha_{d=k}; p_M)}{\varepsilon}\right) - H\left(\frac{\bar{V}_{d=k}}{\varepsilon}\right) - [A_{d=k}(\alpha_{d=k}; p_M) - \bar{V}_{d=k}].$$

Favorable imbalances $k \geq 1$: pure stay-out equilibrium. For any $p_M \in (0, 1)$, any $\alpha_{d=k} \in [0, 1]$, and any $k \geq 1$, the stage-3 ask is at least the prior mean, so

$$A_{d=k}(\alpha_{d=k}; p_M) \geq \bar{V}_{d=k} \geq \bar{V}_{d=1} = \sigma\pi > \varepsilon.$$

Therefore $H(A_{d=k}(\alpha_{d=k}; p_M)/\varepsilon) = H(\bar{V}_{d=k}/\varepsilon) = 1$. Hence

$$\Gamma_{d=k}(\alpha_{d=k}; p_M) = 1 - 1 - [A_{d=k}(\alpha_{d=k}; p_M) - \bar{V}_{d=k}] \leq 0.$$

Thus, for any $p_M \in (0, 1)$ and any $k \geq 1$, $\alpha_{d=k}^* = 0$ is a stable fixed point and is clearly the smallest.

Hopeless imbalances $k \leq -2$: pure stay-out equilibrium. For any $p_M \in (0, 1)$ and any $k \leq -2$,

$$A_{d=k}(0; p_M) = \sigma \frac{r^{k+1} - 1}{r^{k+1} + 1} \leq \sigma \frac{r^{-2+1} - 1}{r^{-2+1} + 1} = \bar{V}_{d=-1} = -\sigma\pi < -\varepsilon.$$

Since the ask strictly decreases in $\alpha_{d=k}$ and is at least the prior mean,

$$\bar{V}_{d=k} \leq A_{d=k}(\alpha_{d=k}; p_M) \leq A_{d=k}(0; p_M) < -\varepsilon.$$

Consequently $H(A_{d=k}(\alpha_{d=k}; p_M) / \varepsilon) = H(\bar{V}_{d=k} / \varepsilon) = 0$, and

$$\Gamma_{d=k}(\alpha_{d=k}; p_M) = 0 - 0 - [A_{d=k}(\alpha_{d=k}; p_M) - \bar{V}_{d=k}] \leq 0.$$

Thus, for any $p_M \in (0, 1)$ and any $k \leq -2$, $\alpha_{d=k}^* = 0$ is a stable fixed point and is clearly the smallest.

Boundary imbalance $k = 0$: pure push equilibrium. $k = 0$ can occur only when n is even. Fix $p_M \in (0, 1)$. Here

$$\bar{V}_{d=0} = 0,$$

and the stage-3 ask is

$$A_{d=0}(\alpha_{d=0}; p_M) = \frac{\sigma\pi(1 - p_M)}{1 - p_M + 2p_M\alpha_{d=0}}.$$

At $\alpha_{d=0} = 0$,

$$A_{d=0}(0; p_M) = \sigma\pi > \varepsilon,$$

so

$$\Gamma_{d=0}(0; p_M) = 1 - \frac{1}{2} - \sigma\pi = \frac{1}{2} - \sigma\pi > 0.$$

Thus $\alpha_{d=0} = 0$ is not a fixed point. For any $\alpha_{d=0} \in [0, 1]$, the ask is

$$A_{d=0}(\alpha_{d=0}; p_M) \in \left[\frac{\sigma\pi(1-p_M)}{1+p_M}, \sigma\pi \right] \subset \left(0, \frac{1}{2} \right).$$

If $A_{d=0}(\alpha_{d=0}; p_M) \geq \varepsilon$, then

$$\Gamma_{d=0}(\alpha_{d=0}; p_M) = \frac{1}{2} - A_{d=0}(\alpha_{d=0}; p_M) \geq \frac{1}{2} - \sigma\pi > 0.$$

If $0 < A_{d=0}(\alpha_{d=0}; p_M) < \varepsilon$, then

$$\Gamma_{d=0}(\alpha_{d=0}; p_M) = A_{d=0}(\alpha_{d=0}; p_M) \left(\frac{1}{2\varepsilon} - 1 \right) > 0$$

because $\varepsilon < 1/2$. Thus no interior fixed point exists either, and the selected continuation equilibrium is the endpoint

$$\alpha_{d=0}^* = 1.$$

Boundary imbalance $k = -1$: mixed or pure push. $k = -1$ can occur only when n is odd. Fix $p_M \in (0, 1)$. Here

$$\bar{V}_{d=-1} = -\sigma\pi,$$

and

$$A_{d=-1}(\alpha_{d=-1}; p_M) = \frac{-\sigma\pi p_M \alpha_{d=-1}}{p_M \alpha_{d=-1} + (1-p_M) \frac{1-\pi^2}{2}}.$$

This is exactly the unfavorable-path continuation problem from Appendix A.5.1. Therefore, with

$$p_0 = \frac{\varepsilon(1-\pi^2)(1-2\sigma\pi)}{2(\sigma\pi - \varepsilon) + \varepsilon(1-\pi^2)(1-2\sigma\pi)},$$

the selected continuation equilibrium is

$$\alpha_{d=-1}^*(p_M) = \begin{cases} 1, & 0 < p_M \leq p_0, \\ \frac{1 - p_M}{p_M} \frac{\varepsilon(1 - \pi^2)(1 - 2\sigma\pi)}{2(\sigma\pi - \varepsilon)}, & p_0 < p_M < 1. \end{cases}$$

The corresponding equilibrium ask is

$$A_{d=-1}^*(p_M) = \begin{cases} -\frac{2\sigma\pi p_M}{(1 - \pi^2) + (1 + \pi^2)p_M}, & 0 < p_M \leq p_0, \\ -\frac{\varepsilon(1 - 2\sigma\pi)}{1 - 2\varepsilon}, & p_0 < p_M < 1. \end{cases}$$

Thus $d = -1$ is the unique pivotal imbalance when n is odd.

A.7.3 Optimal manipulator entry probability

Both the pricing of the binary contract and the manipulator's stage-1 entry problem use the selected continuation equilibrium from Appendix A.7.2. Appendix A.7.1 shows that the stage-2 posterior depends on h_2 only through the imbalance $d(h_2)$, so we can aggregate over imbalances rather than individual price paths. Combining Appendix A.7.1, the general ask and bid formulas (4) and (5), and Appendix A.7.2, we obtain the following stage-3 asks, bids, and no-trade posterior means for a stage-2 price path h_2 with $d(h_2) = k$:

$$A_{d=k}^*(p_M) = \begin{cases} \sigma \frac{r^{k+1} - 1}{r^{k+1} + 1} = \bar{V}_{d=k+1}, & k \geq 1 \text{ or } k \leq -2, \\ \frac{\sigma\pi(1 - p_M)}{1 + p_M}, & k = 0, \\ -\frac{2\sigma\pi p_M}{(1 - \pi^2) + (1 + \pi^2)p_M}, & k = -1 \text{ and } 0 < p_M \leq p_0, \\ -\frac{\varepsilon(1 - 2\sigma\pi)}{1 - 2\varepsilon}, & k = -1 \text{ and } p_0 < p_M < 1, \end{cases}$$

$$B_{d=k} = \sigma \frac{r^{k-1} - 1}{r^{k-1} + 1} = \bar{V}_{d=k-1}.$$

$$\bar{V}_{d=k} = \sigma \frac{r^k - 1}{r^k + 1}.$$

For each stage-2 price path h_2 with $d(h_2) = k$, the expected binary-contract payoffs are

$$\Pi_M^+(p_M, k) = \alpha_{d=k}^*(p_M) H(A_{d=k}^*(p_M)/\varepsilon) + (1 - \alpha_{d=k}^*(p_M)) H(\bar{V}_{d=k}/\varepsilon),$$

$$\Pi_L^+(p_M, k) = q_{d=k,+}(0) H(A_{d=k}^*(p_M)/\varepsilon) + q_{d=k,-}(0) H(B_{d=k}/\varepsilon).$$

Even n . Suppose n is even. Substituting the selected continuation equilibrium and the stage-3 quotes gives

$$\Pi_M^+(p_M, k) = \begin{cases} 1, & k \geq 2, \\ H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right), & k = 0, \\ 0, & k \leq -2, \end{cases}$$

$$\Pi_L^+(p_M, k) = \begin{cases} 1, & k \geq 2, \\ \frac{1}{2} H\left(\frac{\sigma\pi(1-p_M)}{\varepsilon(1+p_M)}\right), & k = 0, \\ 0, & k \leq -2. \end{cases}$$

The probability of imbalance $k = 0$ is $q_{d=0}$. By (10), the positive and negative tails have equal probability, so

$$\sum_{k \geq 2} q_{d=k} = \sum_{k \leq -2} q_{d=k} = \frac{1 - q_{d=0}}{2}.$$

Aggregating over k gives

$$\begin{aligned}\Pi_M^+(p_M) &= \frac{1 - q_{d=0}}{2} + q_{d=0}H\left(\frac{\sigma\pi(1 - p_M)}{\varepsilon(1 + p_M)}\right), \\ \Pi_L^+(p_M) &= \frac{1 - q_{d=0}}{2} + \frac{q_{d=0}}{2}H\left(\frac{\sigma\pi(1 - p_M)}{\varepsilon(1 + p_M)}\right), \\ K^+(p_M) &= q_{d=0}\frac{\sigma\pi(1 - p_M)}{1 + p_M}.\end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}A_1^{P, BE}(p_M) &= p_M \Pi_M^+(p_M) + (1 - p_M) \Pi_L^+(p_M) \\ &= \frac{1 - q_{d=0}}{2} + \frac{q_{d=0}(1 + p_M)}{2}H\left(\frac{\sigma\pi(1 - p_M)}{\varepsilon(1 + p_M)}\right) \\ &= \begin{cases} \frac{1}{2} + \frac{q_{d=0}p_M}{2}, & \frac{\sigma\pi(1 - p_M)}{1 + p_M} \geq \varepsilon, \\ \frac{1}{2} + \frac{q_{d=0}(1 - p_M)(\sigma\pi - \varepsilon)}{4\varepsilon}, & \frac{\sigma\pi(1 - p_M)}{1 + p_M} < \varepsilon. \end{cases}\end{aligned}$$

which is strictly greater than $1/2$ for any $p_M \in (0, 1)$. Therefore the non-crossing equilibrium stage-1 ask is just the break-even stage-1 ask:

$$A_1^{P*}(p_M) = \max\left\{A_1^{P, BE}(p_M), \frac{1}{2}\right\} = A_1^{P, BE}(p_M),$$

and

$$B_1^{P*}(p_M) = 1 - A_1^{P*}(p_M)$$

for any $p_M \in (0, 1)$.

The manipulator's expected profit conditional on entering at stage 1 is

$$\begin{aligned}
 & S_M^+(p_M) \\
 &= \Pi_M^+(p_M) - K^+(p_M) - A_1^{P*}(p_M) \\
 &= \frac{q_{d=0}}{2} \left[(1 - p_M) H \left(\frac{\sigma\pi(1 - p_M)}{\varepsilon(1 + p_M)} \right) - \frac{2\sigma\pi(1 - p_M)}{1 + p_M} \right].
 \end{aligned}$$

Hence

$$\bar{S}_M^+(p_M) = p_M S_M^+(p_M) = \frac{q_{d=0}}{2} \Phi(p_M; \sigma, \pi, \varepsilon),$$

where Φ is exactly the same function as defined in (8) for the manipulator's optimization problem under the two-ordinary-round extension. Thus every even n has the same optimal entry probability as the two-ordinary-round extension:

$$p_M^{*(n)} = p_M^{*(2)} \quad \text{if } n \text{ is even.}$$

Restoring horizon superscripts, since $q_{d=0} = m^{(n)}$ for even n and $m^{(2)} = (1 - \pi^2)/2$, where $m^{(n)}$ is defined in Section 5.4 and $m^{(2)}$ in Section 5.4, the preceding formulas imply

$$\begin{aligned}
 S_M^{+(n)}(p_M^{*(n)}) &= \frac{m^{(n)}}{m^{(2)}} S_M^{+(2)}(p_M^{*(2)}), \\
 \bar{S}_M^{+(n)}(p_M^{*(n)}) &= \frac{m^{(n)}}{m^{(2)}} \bar{S}_M^{+(2)}(p_M^{*(2)})
 \end{aligned}$$

for even n .

Odd n . Suppose n is odd. Substituting the selected continuation equilibrium and the stage-3 quotes gives

$$\Pi_M^+(p_M, k) = \begin{cases} 1, & k \geq 1, \\ \alpha_{d=-1}^*(p_M) H(A_{d=-1}^*(p_M)/\varepsilon), & k = -1, \\ 0, & k \leq -3, \end{cases}$$

$$\Pi_L^+(p_M, k) = \begin{cases} 1, & k \geq 3, \\ \frac{3 + \pi^2}{4}, & k = 1, \\ \frac{1 - \pi^2}{2} H(A_{d=-1}^*(p_M)/\varepsilon), & k = -1, \\ 0, & k \leq -3. \end{cases}$$

The imbalances $k = \pm 1$ have probabilities $q_{d=\pm 1}$. By (10), $q_{d=1} = q_{d=-1}$ and the positive and negative tails have equal probability, so

$$\sum_{k \geq 3} q_{d=k} = \sum_{k \leq -3} q_{d=k} = \frac{1}{2} - q_{d=-1}.$$

Aggregating over k gives

$$\begin{aligned} \Pi_M^+(p_M) &= \frac{1}{2} + q_{d=-1} \alpha_{d=-1}^*(p_M) H(A_{d=-1}^*(p_M)/\varepsilon), \\ \Pi_L^+(p_M) &= \frac{1}{2} + q_{d=-1} (1 - \pi^2) \left[\frac{1}{2} H(A_{d=-1}^*(p_M)/\varepsilon) - \frac{1}{4} \right], \\ K^+(p_M) &= q_{d=-1} \alpha_{d=-1}^*(p_M) [A_{d=-1}^*(p_M) + \sigma\pi]. \end{aligned}$$

The unconstrained shadow stage-1 ask obtained from the prediction-market maker's break-even condition is

$$\begin{aligned}
 A_1^{P, BE}(p_M) &= p_M \Pi_M^+(p_M) + (1 - p_M) \Pi_L^+(p_M) \\
 &= \frac{1}{2} + q_{d=-1} \left[p_M \alpha_{d=-1}^*(p_M) H(A_{d=-1}^*(p_M) / \varepsilon) + (1 - p_M) (1 - \pi^2) \left[\frac{1}{2} H(A_{d=-1}^*(p_M) / \varepsilon) - \frac{1}{4} \right] \right] \\
 &= \begin{cases} \frac{1}{2} - q_{d=-1} \frac{p_M (\sigma\pi - \varepsilon)}{2\varepsilon}, & 0 < p_M \leq p_0, \\ \frac{1}{2} - q_{d=-1} \frac{(1 - p_M) (1 - \pi^2) (1 - 2\sigma\pi)}{4}, & p_0 < p_M < 1. \end{cases}
 \end{aligned}$$

which is strictly less than $1/2$ for any $p_M \in (0, 1)$. Therefore the non-crossing equilibrium stage-1 ask and bid are respectively

$$A_1^{P*}(p_M) = \max \left\{ A_1^{P, BE}(p_M), \frac{1}{2} \right\} = \frac{1}{2}, \quad B_1^{P*}(p_M) = 1 - A_1^{P*}(p_M) = \frac{1}{2}$$

for any $p_M \in (0, 1)$.

The manipulator's expected profit conditional on entering at stage 1 is

$$S_M^+(p_M) = \Pi_M^+(p_M) - K^+(p_M) - \frac{1}{2},$$

which yields

$$S_M^+(p_M) = \begin{cases} q_{d=-1} \frac{(1 - \pi^2) (1 - 2\sigma\pi) (1 - p_M/p_0)}{2 [(1 - \pi^2) + (1 + \pi^2) p_M]}, & 0 < p_M \leq p_0, \\ 0, & p_0 < p_M < 1. \end{cases}$$

The pure-push branch is positive when $0 < p_M < p_0$ and vanishes at $p_M = p_0$. Therefore, the

manipulator maximizes the ex-ante profit

$$\bar{S}_M^+(p_M) = p_M S_M^+(p_M) = q_{d=-1} \frac{p_M (1 - \pi^2) (1 - 2\sigma\pi) (1 - p_M/p_0)}{2 [(1 - \pi^2) + (1 + \pi^2) p_M]}$$

over $0 < p_M < p_0$. This equals $2q_{d=-1}$ times the manipulator's ex-ante profit (7) in the baseline model. Thus every odd n has the same optimal entry probability as the baseline model:

$$p_M^{*(n)} = p_M^{*(1)} \quad \text{if } n \text{ is odd.}$$

Since $p_M^{*(n)} = p_M^{*(1)} \in (0, p_0)$, the optimum lies on the pure-push branch. Restoring horizon superscripts, since $q_{d=-1} = m^{(n)}$ for odd n and $m^{(1)} = 1/2$, where $m^{(n)}$ is defined in Section 5.4 and $m^{(1)}$ in Section 5.4, the preceding formulas imply

$$\begin{aligned} S_M^{+(n)}(p_M^{*(n)}) &= \frac{m^{(n)}}{m^{(1)}} S_M^{+(1)}(p_M^{*(1)}), \\ \bar{S}_M^{+(n)}(p_M^{*(n)}) &= \frac{m^{(n)}}{m^{(1)}} \bar{S}_M^{+(1)}(p_M^{*(1)}) \end{aligned}$$

for odd n . The same observation pins down

$$\alpha_{d=-1}^{*(n)}(p_M^{*(n)}) = 1 \quad \text{if } n \text{ is odd.}$$

A.7.4 Price discovery

Fix $n > 2$. In this subsection only, write $q_{d=k}^{(n)}$ for the imbalance- k probability in (9) computed under the n -ordinary-round extension to emphasize the dependence on n . For any stage-2 price path h_2 with $d(h_2) = k$ and any $p_M^{(n)} \in (0, 1)$, also write $q_{d=k, h_3}(p_M^{(n)})$, $\bar{V}_{d=k, h_3}(p_M^{(n)})$, and $v_{d=k, h_3}(p_M^{(n)})$ for the corresponding conditional stage-3 outcome probability, posterior mean, and posterior variance.

Then

$$\mathbb{E}[v_{h_2 h_3}^{(n)}](p_M^{(n)}) = \sum_{k \in \mathcal{D}_n} q_{d=k}^{(n)} \sum_{h_3 \in \{+, \emptyset, -\}} q_{d=k, h_3}(p_M^{(n)}) v_{d=k, h_3}(p_M^{(n)})$$

by the law of total expectation. To prove the price-discovery part of Proposition 7, it suffices to compare horizon n to horizon $n - 1$ through the no-manipulation $(n - 1)$ -stage benchmark. We show, for every $p_M^{(n)}, p_M^{(n-1)} \in (0, 1)$,

$$\mathbb{E} \left[v_{h_2 h_3}^{(n)} \right] (p_M^{(n)}) < \mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (0) < \mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (p_M^{(n-1)}).$$

At the no-manipulation benchmark $p_M^{(n-1)} = 0$, the $(n - 1)$ -ordinary-round model has $n - 1$ ordinary spot rounds in stage 2 and one ordinary spot round in stage 3, and thus stages 2 and 3 together share the same event structure as stage 2 under the n -ordinary-round extension. Hence

$$\mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (0) = \sum_{k \in \mathcal{D}_n} q_{d=k}^{(n)} v_{d=k}.$$

For each $k \in \mathcal{D}_n$, the law of total variance gives

$$\sum_{h_3 \in \{+, \emptyset, -\}} q_{d=k, h_3}^{(n)} (p_M^{(n)}) v_{d=k, h_3}^{(n)} (p_M^{(n)}) = v_{d=k} - \text{Var} \left(\bar{V}_{d=k, h_3} \mid d(h_2) = k \right) (p_M^{(n)}).$$

Since $\rho_{d=k} \in (0, 1)$,

$$q_{d=k, -}^{(n)} (p_M^{(n)}) = \left(1 - p_M^{(n)} \right) \left[\frac{1 - \pi}{2} + \pi (1 - \rho_{d=k}) \right] > 0,$$

$$\bar{V}_{d=k, -}^{(n)} (p_M^{(n)}) - \bar{V}_{d=k} = B_{d=k} - \bar{V}_{d=k} = -\frac{\sigma \pi [1 - (2\rho_{d=k} - 1)^2]}{1 - (2\rho_{d=k} - 1)\pi} < 0.$$

Therefore

$$\text{Var} \left(\bar{V}_{d=k, h_3} \mid d(h_2) = k \right) (p_M^{(n)}) \geq q_{d=k, -}^{(n)} (p_M^{(n)}) \left[\bar{V}_{d=k, -}^{(n)} (p_M^{(n)}) - \bar{V}_{d=k} \right]^2 > 0.$$

Taking the $q_{d=k}^{(n)}$ -weighted average of the law-of-total-variance identity over $k \in \mathcal{D}_n$ gives

$$\mathbb{E} \left[v_{h_2 h_3}^{(n)} \right] (p_M^{(n)}) = \sum_{k \in \mathcal{D}_n} q_{d=k}^{(n)} \left[v_{d=k} - \text{Var} \left(\bar{V}_{d=k, h_3} \mid d(h_2) = k \right) (p_M^{(n)}) \right].$$

Each $q_{d=k}^{(n)}$ is strictly positive, and the previous display shows that each subtracted posterior-mean dispersion term is strictly positive, so

$$\mathbb{E} \left[v_{h_2 h_3}^{(n)} \right] (p_M^{(n)}) < \sum_{k \in \mathcal{D}_n} q_{d=k}^{(n)} v_{d=k} = \mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (0).$$

Finally, although Proposition 3 is stated for the baseline model, its proof in Appendix A.4 conditions only on the stage-2 posterior and thus still applies to the $(n-1)$ -ordinary-round extension, giving, for every $p_M^{(n-1)} \in (0, 1)$,

$$\mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (p_M^{(n-1)}) > \mathbb{E} \left[v_{h_2 h_3}^{(n-1)} \right] (0).$$


Chaining the inequalities gives the stated result.

B Polymarket Contract Trading Volume


This appendix documents the trading volume of the Polymarket crypto up/down contracts that the main analysis studies, across the six coins (BTC, ETH, SOL, XRP, DOGE, BNB) and the main contract durations (five-minute, fifteen-minute, four-hour). The contracts carry tens of millions of dollars of daily volume; the five-minute contract is the most heavily traded and BTC dominates the cross-section.

By coin. Figure A1a stacks daily Polymarket crypto up/down USDC volume by coin, summed across all contract durations, from the earliest Polymarket trade. Volume grows from a few million dollars a day in October 2025 to roughly \$30–40M by March 2026, with a visible step up when the five-minute contract launches in February. BTC is the large majority of the total throughout, ETH a distant second, and SOL/XRP/DOGE/BNB thin slivers, consistent with the manipulation footprint being concentrated in the coin whose pool is deepest.

By contract duration. Figure A1b stacks daily BTC volume by contract duration. The fifteen-minute contract carries most of the volume through the pre-five-minute period and the four-hour contract a thin sliver; the five-minute layer appears at its February launch and quickly becomes the largest, raising total daily BTC volume to roughly \$30M.



(a) All coins, by coin.



(b) BTC, by contract duration.

Figure A1. Daily Polymarket crypto up/down USDC volume. This figure plots stacked daily taker USDC notional on the Polymarket crypto up/down contracts, from the earliest Polymarket trade (the October 2025 fifteen-minute / four-hour launch). (a) All six coins, summed across the three durations (five-minute, fifteen-minute, four-hour), stacked by coin. (b) BTC only, stacked by contract duration (four-hour, fifteen-minute, five-minute, from bottom to top).

C Overlap and Moneyness Refinements

The two subsections below refine the cycle-level partition under which the main-body activity and reversal results are estimated. The first refines the within-P3 activity regression (Table 4) and reversal regression (Table 5) into the full 15-min-overlap $\times$ per-contract-NTM groups. The second applies the moneyness cut to the fifteen-minute contract alone in P2, before the five-minute launch, asking whether the NTM concentration already appears at the longer horizon.

Proof of Proposition 8. Since $1 - \pi^2 \in (0, 1)$ and $\binom{n}{\lfloor n/2 \rfloor} \leq 2^n$,

$$0 \leq m^{(n)} \leq (1 - \pi^2)^{\lfloor n/2 \rfloor}.$$

The limit of $m^{(n)}$ then follows from the squeeze theorem:

$$\lim_{n \uparrow \infty} (1 - \pi^2)^{\lfloor n/2 \rfloor} = 0.$$

Appendix A.7.3 shows that the optimal profit objects are the corresponding lower-horizon objects scaled by the pivotal probabilities:

$$S_M^{+(n)}(p_M^{*(n)}) = \begin{cases} \frac{m^{(n)}}{m^{(1)}} S_M^{+(1)}(p_M^{*(1)}), & n \text{ is odd,} \\ \frac{m^{(n)}}{m^{(2)}} S_M^{+(2)}(p_M^{*(2)}), & n \text{ is even,} \end{cases}$$

and similarly

$$\bar{S}_M^{+(n)}(p_M^{*(n)}) = \begin{cases} \frac{m^{(n)}}{m^{(1)}} \bar{S}_M^{+(1)}(p_M^{*(1)}), & n \text{ is odd,} \\ \frac{m^{(n)}}{m^{(2)}} \bar{S}_M^{+(2)}(p_M^{*(2)}), & n \text{ is even.} \end{cases}$$

Clearly, as $n \uparrow \infty$, both profit objects converge to zero along with $m^{(n)}$. □

C.1 Activity and reversal by overlap and moneyness

This subsection refines the within-P3 NTM results, the activity regression of Table 4 and the reversal regression of Table 5, by whether the cycle's close also coincides with a 15-minute mark, where the 15m Polymarket contract settles alongside the 5m contract. The cycle groups are written in set notation: $5m \cap 15m$ denotes overlapping cycles, where both the 5m and 15m Polymarket contracts settle at the close (four of twelve per hour), and $5m \setminus 15m$ non-overlapping cycles, where only the 5m contract settles. NTM is per contract from the Polymarket Up-token price near time-to-maturity $TTM \in [10, 12]$ s: NTM_{5m}^{poly} iff $|p_{Up,5m} - 0.5| < 0.10$, NTM_{15m}^{poly} iff $|p_{Up,15m} - 0.5| < 0.10$ (defined only for overlapping cycles). Together these yield six mutually exclusive groups: $5m \setminus 15m$ split by 5m-contract NTM, and $5m \cap 15m$ split into neither-NTM, NTM_{5m} -only, NTM_{15m} -only, and both-NTM.

Activity. Restricting to P3 and to bins in $\{0-24, 29\}$, the regression is

$$\begin{aligned} \log y_{c,b} = & \alpha_d + \gamma_b + \mu_h + \delta_{5m \cap 15m} \mathbf{1}\{b=29\} \cdot D_{5m \cap 15m,c} \\ & + \delta_{NTM_{5m}, 5m \setminus 15m} \mathbf{1}\{b=29\} \cdot D_{5m \setminus 15m,c} \cdot \mathbf{1}\{NTM_{5m,c}^{poly}\} \\ & + \delta_{NTM_{5m}, 5m \cap 15m} \mathbf{1}\{b=29\} \cdot D_{5m \cap 15m,c} \cdot \mathbf{1}\{NTM_{5m,c}^{poly}\} \\ & + \delta_{NTM_{15m}, 5m \cap 15m} \mathbf{1}\{b=29\} \cdot D_{5m \cap 15m,c} \cdot \mathbf{1}\{NTM_{15m,c}^{poly}\} \\ & + \delta_{NTM_{5m} \times NTM_{15m}, 5m \cap 15m} \mathbf{1}\{b=29\} \cdot D_{5m \cap 15m,c} \cdot \mathbf{1}\{NTM_{5m,c}^{poly}\} \cdot \mathbf{1}\{NTM_{15m,c}^{poly}\} \\ & + (\text{level effects}) + \varepsilon_{c,b}, \end{aligned}$$

estimated with date, hour-of-day, and bin fixed effects and cycle-clustered standard errors. The baseline, absorbed into the bin fixed effect γ_{29} , is the $5m \setminus 15m \times \text{FFM}$ group; each δ is the bin-29-vs-body increase in its group over that baseline, $\delta_{5m \cap 15m}$ in overlapping cycles where neither contract is NTM, $\delta_{NTM_{5m}, 5m \setminus 15m}$ the 5m-contract NTM effect in non-overlapping cycles, $\delta_{NTM_{5m}, 5m \cap 15m}$ the 5m-contract NTM effect in overlapping cycles, and $\delta_{NTM_{15m}, 5m \cap 15m}$ the incremental 15m-contract NTM effect in overlapping cycles. The final term, $\delta_{NTM_{5m} \times NTM_{15m}, 5m \cap 15m}$, is the both-NTM interaction:

it saturates the four overlap-moneyness cells, so the both-NTM cell is a free parameter rather than the additive sum of the two single-contract effects.

In the order-flow column of Table A1, the overlap-without-NTM effect is small ($\hat{\delta}_{5m \cap 15m} = 0.21$), while the 5m-contract NTM effect is large in both the non-overlap group ($\hat{\delta}_{NTM_{5m, 5m \setminus 15m}} = 1.32$) and the overlap group ($\hat{\delta}_{NTM_{5m, 5m \cap 15m}} = 1.39$), and the 15m-contract NTM coefficient adds a further $\hat{\delta}_{NTM_{15m, 5m \cap 15m}} = 1.12$ inside the overlap group. The both-NTM interaction is positive on top of that: in cycles NTM by *both* contracts, the bin-29 spike exceeds the sum of the two single-contract effects by a further $\hat{\delta}_{NTM_{5m \times NTM_{15m}}} = 1.14$ log points (significant at the 5% level despite resting on only 11 such cycles), so the both-NTM overlap close sums all four terms, $0.21 + 1.39 + 1.12 + 1.14 = 3.87$ log points ($e^{3.87} \approx 48\times$ on the typical cycle) above the $5m \setminus 15m$ FFM baseline. This is what the model predicts: with two pools winnable at one close, a push is worth more than the sum of its parts. The same ordering holds for absolute return, and in the Polymarket columns the depth coefficients are again negative (prediction-market makers withdraw). The activity surge thus compounds with both moneyness and contract overlap: it is largest exactly where two contracts settle at one close and the outcome is still live by both.

Reversal. The reversal counterpart estimates, at each lag, a baseline slope plus four additive increments over the same six groups: the $5m \setminus 15m \times$ FFM group is the baseline slope γ_j on r_{29} , and we add $\delta_{5m \cap 15m}$ (overlapping close, neither contract NTM), $\delta_{NTM_{5m, 5m \setminus 15m}}$ (5m-contract NTM in non-overlapping cycles), $\delta_{NTM_{5m, 5m \cap 15m}}$ (5m-contract NTM in overlapping cycles), and $\delta_{NTM_{15m, 5m \cap 15m}}$ (incremental 15m-contract NTM in overlapping cycles), each as a slope increment on r_{29} . The nine overlapping cycles NTM by both contracts load on both overlap-NTM increments; unlike the saturated activity regression, nine cycles cannot support a free interaction slope.

Figure A2 plots these slopes by group. At lag $j = +1$ in Table A2, the moneyness increments carry the result. Against the non-overlapping FFM baseline ($\gamma_{+1} = -0.08$), a non-overlapping cycle that is NTM_{5m} adds a further -0.16 , about tripling the baseline reversal, and within overlapping cycles, NTM by either contract adds about -0.22 ($\delta_{NTM_{5m, 5m \cap 15m}} = -0.23$ and $\delta_{NTM_{15m, 5m \cap 15m}} =$

Table A1. Within-P3 bin-29 activity increase by 15-min overlap and per-contract Polymarket NTM, Binance and Polymarket. This table reports cycle-level OLS estimates on the (cycle, bin) panel restricted to bins {0–24, 29} and P3 cycles, with date, hour-of-day, and bin fixed effects; dependent variable $\log y_{c,b}$, the log of the column metric. Coefficients are log differences (δ is a multiplicative e^δ effect on the typical geometric-mean cycle). Set notation for the cycle group: $5m \cap 15m = 5m$ cycles whose close coincides with a 15-min boundary (both the 5m and 15m contracts settle); $5m \setminus 15m =$ the other eight of twelve per hour (only the 5m contract settles). NTM is per contract from the Polymarket Up-token price near $TTM \in [10, 12]$ s: NTM_{5m}^{poly} iff $|p_{Up,5m} - 0.5| < 0.10$, NTM_{15m}^{poly} iff $|p_{Up,15m} - 0.5| < 0.10$ (overlapping cycles only). The five reported coefficients are $1\{b=29\}$ interactions; the $5m \setminus 15m \times FFM$ bin-29 increase is the baseline, absorbed in the bin FE. The $\delta_{NTM_{5m} \times NTM_{15m}}$ interaction saturates the four overlap-moneyness cells, so each cell's bin-29 increase is a free parameter rather than an additive constraint; on the Binance order-flow sample the cells hold 586 ($5m \setminus 15m \times NTM_{5m}$), 211 (NTM_{5m} -only), 132 (NTM_{15m} -only), and 11 (both-NTM) cycles, so the interaction is estimated from few cycles. One observation per (cycle, bin) cell. HC1 standard errors in parentheses are clustered at the cycle level. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Binance			Polymarket		
	Order flow	Abs. return	Depth	Order flow	Abs. price change	Depth
$\delta_{5m \cap 15m}$	0.212*** (0.043)	0.109*** (0.034)	0.013* (0.008)	0.163*** (0.032)	0.768*** (0.153)	0.026 (0.047)
$\delta_{NTM_{5m}, 5m \setminus 15m}$	1.324*** (0.105)	0.555*** (0.063)	0.154*** (0.026)	−0.090* (0.052)	2.896*** (0.110)	−0.980*** (0.064)
$\delta_{NTM_{5m}, 5m \cap 15m}$	1.393*** (0.185)	0.479*** (0.111)	0.168*** (0.041)	−0.307*** (0.099)	1.879*** (0.176)	−1.146*** (0.113)
$\delta_{NTM_{15m}, 5m \cap 15m}$	1.123*** (0.228)	0.435*** (0.145)	0.205*** (0.062)	0.167 (0.134)	0.251 (0.517)	−0.234 (0.211)
$\delta_{NTM_{5m} \times NTM_{15m}, 5m \cap 15m}$	1.143** (0.583)	0.457* (0.277)	0.259 (0.262)	0.290 (0.341)	−1.331 (1.097)	−0.060 (0.324)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Bin FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	352,370	272,785	352,370	351,791	122,958	126,552
Clusters (cycle)	13,553	13,551	13,553	13,553	5,001	5,001

−0.22), all significant; the deepest-reverting groups, overlap $\times$ NTM, revert about five times as hard as the baseline. That the NTM_{5m} and NTM_{15m} increments are almost identical says each contract's moneyness is on its own enough to mark a high-reversal cycle, consistent with the manipulator-selection reading: the payoff per dollar of push is larger when two contracts settle at the same close. The overlap level increment ($\delta_{5m \cap 15m} = -0.08$) we treat as a descriptive control rather than a settlement effect: 15-minute marks attract round-time algorithmic flow of their own (Section 6.3.3). The manipulation reading therefore rests on the moneyness increments, which are defined by the live contracts' own prices and have no clock analogue.




Figure A2. Within-P3 post-settlement reversal slope γ_j , by cycle group. This figure plots γ_j for each cycle group (the underlying regression is Table A2): left panel non-overlapping ($5m \setminus 15m$) cycles, right panel overlapping ($5m \cap 15m$) cycles, each split by per-contract Polymarket NTM. Markers are γ_j ; error bars are $\pm 1.96 \times$ cluster-robust SE (clustered by date); $j = 0$ is omitted. The both-NTM overlap group (nine cycles) is omitted.

C.2 Moneyness for the fifteen-minute contract alone

Section 6.3.3 finds only a strongly attenuated fifteen-minute footprint, but that test is unconditional: it averages over all overlapping closes, most of which are far from the money by the time the contract settles. Within P3, the footprint concentrates in NTM cycles (Section 6.3.2), so a genuine fifteen-minute effect could hide in the unconditional average, diluted by the FFM majority. This appendix therefore applies the moneyness cut to the fifteen-minute contract on its own. We take P2, where the fifteen-minute (and four-hour) contract is live but the five-minute contract is not, restrict to overlapping cycles, those whose close lands on a 15-min mark, where the fifteen-minute contract settles, and classify them NTM vs FFM by the *fifteen-minute* contract's Polymarket Up-token price within 0.10 of 0.5 near $TTM \in [10, 12]$ s, the same moneyness rule applied to the contract that is actually live. If the fifteen-minute horizon supported manipulation, its footprint should surface here, in exactly the cycles where a push could flip its resolution.

Table A3 reports the result, and it is largely a null. Panel A shows a modest bin-29 activity increase in NTM cycles — about 0.24 in order flow and 0.30 in absolute return, roughly 27% and

Table A2. Within-P3 post-settlement reversal slope by 15-min overlap and per-contract moneyiness. This table reports, at each lag $j \in \{-3, -2, -1, +1, +2, +3\}$, the coefficients from a single cycle-level OLS on P3 cycles (Feb 12 – Apr 8, 2026): the $5m \setminus 15m \times \text{FFM}$ group is the baseline reversal slope γ_j on r_{29} , and the four reported δ are additive slope increments over it, $\delta_{5m \cap 15m}$ (overlapping close), $\delta_{\text{NTM}_{5m}, 5m \setminus 15m}$ and $\delta_{\text{NTM}_{5m}, 5m \cap 15m}$ (5m-contract NTM in non-overlapping and overlapping cycles), and $\delta_{\text{NTM}_{15m}, 5m \cap 15m}$ (incremental 15m-contract NTM in overlapping cycles); the stars test each coefficient against zero. Groups use set notation: $5m \setminus 15m$ = non-overlapping cycles (only the 5m contract settles), $5m \cap 15m$ = overlapping cycles whose close coincides with a 15-min mark (the 5m and 15m contracts settle simultaneously). Within each, cycles are split by per-contract Polymarket NTM, defined by the Up-token price near TTM $\in [10, 12]$ s: $\text{NTM}_{5m}^{\text{poly}}$ iff $|p_{Up, 5m} - 0.5| < 0.10$, $\text{NTM}_{15m}^{\text{poly}}$ iff $|p_{Up, 15m} - 0.5| < 0.10$. The nine overlapping cycles NTM by *both* contracts load on both the NTM_{5m} and NTM_{15m} overlap increments — their slope is constrained to the additive sum (unlike the saturated activity Table A1; nine cycles cannot support a free interaction slope). Date, hour-of-day, and cycle-group fixed effects; standard errors in parentheses are clustered by date. Sample is P3 cycles with the relevant Polymarket classifications available: $5m \setminus 15m$ cycles need the 5m classifier, $5m \cap 15m$ cycles need both. Negative j = same-cycle bin ($29 + j$); positive j = bin ($j - 1$) of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
$5m \setminus 15m$, FFM (baseline)	0.001 (0.015)	0.031 (0.019)	0.052*** (0.015)	-0.083*** (0.028)	-0.028 (0.022)	0.019 (0.014)
$\delta_{5m \cap 15m}$	-0.027 (0.022)	-0.023 (0.025)	-0.026 (0.021)	-0.078** (0.033)	-0.039 (0.037)	-0.011 (0.028)
$\delta_{\text{NTM}_{5m}, 5m \setminus 15m}$	-0.020 (0.022)	-0.019 (0.032)	0.007 (0.028)	-0.163*** (0.044)	-0.003 (0.039)	-0.053 (0.039)
$\delta_{\text{NTM}_{5m}, 5m \cap 15m}$	0.035 (0.032)	-0.012 (0.032)	-0.072** (0.035)	-0.225** (0.097)	0.062 (0.069)	0.035 (0.040)
$\delta_{\text{NTM}_{15m}, 5m \cap 15m}$	0.068 (0.045)	-0.020 (0.043)	0.009 (0.054)	-0.216** (0.106)	0.115 (0.071)	-0.064 (0.072)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Cycle-group FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,507					
Clusters (UTC date)	56					

35% more close-time activity — both significant at the 5% level, but no effect on depth. These magnitudes are a small fraction of the five-minute effect in P3: differencing the log coefficients, the P3 NTM order-flow effect of Table 4 is about $e^{1.37-0.24} \approx 3.1\times$ larger, and P3 carries the depth and reversal signatures that P2 lacks. Panel B shows the reversal channel is absent altogether: the FFM baseline reverts at the first post-close lag (the general overlap reversal, present once the fifteen-minute contract is live), but the NTM increment $\delta_{\text{NTM}, +1}$ is small, wrong-signed, and far from significant. The thin NTM sample (about 230 cycles) is the binding constraint on Panel B. Taken together, the fifteen-minute contract alone generates at most a modest moneyiness-graded

Table A3. Fifteen-minute-contract moneyness in P2 (BTC). This table reports near-settlement activity (Panel A) and post-settlement reversal (Panel B) in P2 overlapping cycles, split by fifteen-minute-contract Polymarket NTM. Sample: P2 cycles (Oct 9, 2025 – Feb 11, 2026; fifteen-minute and four-hour contracts live, five-minute not yet) whose close coincides with a 15-min mark, classified NTM iff the fifteen-minute-contract Up-token price is within 0.10 of 0.5 near $TTM \in [10, 12]$ s, FFM otherwise. *Panel A:* the bin-29 NTM activity increase δ_{NTM} (the coefficient on $\mathbf{1}\{b=29\} \cdot \mathbf{1}\{NTM\}$) for each metric, log-OLS on the (cycle, bin) panel restricted to bins $\{0-24, 29\}$ with date, hour-of-day, and bin fixed effects; standard errors clustered at the cycle level. Metrics as in Table 4; coefficients are log differences, so δ is a multiplicative e^δ effect. *Panel B:* the post-settlement reversal slope by lag j , the FFM baseline slope γ_j on r_{29} and the NTM increment $\delta_{NTM,j}$, with date, hour-of-day, and group fixed effects; standard errors clustered by date; a negative slope is reversal. Negative j = same-cycle bin $(29 + j)$; positive j = bin $(j - 1)$ of the next cycle. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

<i>Panel A: Near-settlement activity</i>						
	Order flow		Abs. return	Depth		
δ_{NTM}	0.243**		0.303**	-0.015		
	(0.116)		(0.144)	(0.017)		
Date FE	Yes		Yes	Yes		
Hour-of-day FE	Yes		Yes	Yes		
Bin FE	Yes		Yes	Yes		
Observations	158,570		102,329	158,570		
NTM cycles			229			
FFM cycles			5,870			

<i>Panel B: Post-settlement reversal</i>						
	Pre-close			Post-close		
	$j = -3$	$j = -2$	$j = -1$	$j = +1$	$j = +2$	$j = +3$
FFM γ_j	0.121	-0.026	0.015	-0.137***	-0.011	-0.029
	(0.108)	(0.036)	(0.028)	(0.045)	(0.051)	(0.037)
$\delta_{NTM,j}$	-0.082	-0.129*	-0.004	0.076	-0.184	0.011
	(0.115)	(0.077)	(0.061)	(0.234)	(0.115)	(0.071)
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Hour-of-day FE	Yes	Yes	Yes	Yes	Yes	Yes
Cycle-group FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations				6,069		
Clusters (UTC date)				125		

activity increase, with neither the liquidity withdrawal-versus-provision pattern nor the deeper NTM reversal that define the P3 footprint. Even in the cycles where a push could flip the fifteen-minute resolution, the footprint barely appears: the fifteen-minute window is too long for the push to be reliably pivotal, and the near-settlement effect switches on with the five-minute contract.

Table A4. Binance spot midquote vs. the Chainlink resolution price, BTC. This table compares the Binance midquote with the Chainlink price that settles each market, one observation per Polymarket BTC market with both the Chainlink strike (fixed at window open) and the Chainlink settlement print (at window close) known. *Median Basis_t* is the median per-market (Chainlink – Binance mid)/Binance mid at window open (bps). *SD ΔBasis_t* is the standard deviation of the open-to-close change $\Delta Basis_t \equiv Basis_{close} - Basis_{open}$ (bps). *Strike-side agreement* is the fraction of markets on which the Binance mid finishes on the same side of the strike as the realized resolution, $\text{sign}(\text{mid}_{close} - \text{strike}) = \text{sign}(\text{settlement} - \text{strike})$.

Interval	<i>n</i>	Median <i>Basis_t</i>	SD $\Delta Basis_t$	Strike-side agreement
5m	5,553	–2.45	1.44	84.8%
15m	1,855	–2.49	2.24	91.3%

D Binance Spot Tracks the Chainlink Resolution Price

The 5m / 15m / 4h crypto up/down contracts resolve on a Chainlink VWAP price, not on Binance. Every Binance-based result in the main body therefore rests on the Binance midquote being a faithful proxy for the price that actually settles each market. We check this directly on the subset of BTC markets for which the Chainlink strike (fixed at window open) and the Chainlink settlement print (at window close) are both known; coverage is concentrated in the Mar 16 – Apr 10, 2026 window. For each such market we read the Binance midquote at the window open and close from the 10-second midquote series via a strictly causal match: each bucket stores the *last* book mid in $[t, t + 10 \text{ s})$, so its effective timestamp is the bucket end, and we match each window boundary to the last book mid of the 10-second bucket ending at or before that instant.

Define, per market, the Chainlink–Binance basis at time t (with mid_t the Binance midquote):

$$Basis_t = (\text{Chainlink}_t - \text{mid}_t) / \text{mid}_t.$$

The market’s realized Up/Down outcome is $\text{sign}(\text{settlement} - \text{strike})$; the directional call a manipulator targets is whether the Binance mid finishes above or below the strike, $\text{sign}(\text{mid}_{close} - \text{strike})$, the side a final-10s push tries to land on.

Table A4 reports the comparison. The basis is a stable level offset of ≈ -2.5 bps and is essentially flat across horizons, so it cancels in the open-to-close return comparison: the within-window change

in the basis has standard deviation of only $\approx 1.4\text{--}2.2$ bps, small relative to a typical window move (5.8 bps at 5m, 10.8 bps at 15m), so within each window the Binance mid and the Chainlink price move essentially together.

Where Binance and Chainlink *can* part company is right at the strike. Because Binance sits ≈ 2.5 bps above Chainlink, the Binance mid can finish just above the strike while Chainlink, and hence the resolution, is still below it. The strike-side agreement in the last column is therefore below one: 84.8% at 5m and 91.3% at 15m, rising with horizon as the typical move grows relative to the fixed ≈ 2.5 bps basis. The residual disagreement is confined to genuine coin-toss windows: sorting the BTC 5m markets into quintiles by the size of the Binance window move, agreement is 61% in the smallest quintile ($|\text{move}| < 1.8$ bps), 87% in the middle, and 100% once the move exceeds ≈ 13 bps. That ≈ 2.5 bps band is precisely the margin a spot push must clear: to flip the resolution a manipulator must push Binance *past* the strike by more than the basis, not merely to it; a push that only reaches the strike still fails on the $\sim 15\%$ of 5m windows where the basis leaves Chainlink on the other side. A push well past the strike, however, reliably carries the resolution, which makes spot-market pressure a coherent channel for moving the prediction-market outcome.

E Manipulator Cohort Threshold Sensitivity

The manipulator cohort definition in Section 6.4.4 uses two ad hoc thresholds: a minimum manipulated-cycle participation count of 5 and a minimum aggregate manipulated-cycle PnL of \$2,000. Table A5 re-derives the *full* trader-type attribution at each cell of a 4×5 grid of $(n_{\text{cycle},\min}, \$_{\min})$ thresholds: the manipulator cohort is recomputed, the MM cohort is re-derived as the behavioral residual (same passive-rate, fills-per-active-cycle, and inventory criteria as the main text), and retail/other is the final residual, so the entire partition responds to the threshold. For each cell the three panels report the three cohort metrics of Table 10: net PnL (Panel A), the fraction of cycles in which each cohort is net negative (Panel B), and the pool-weighted loss share (Panel C), each split across manipulated and normal cycles. Within each regime the three cohort PnLs sum to zero and the three loss shares to 100% by construction.

All three metrics are stable across the twenty cells. The manipulated-cycle wealth transfer runs directly from retail to manipulators: manipulators net between +\$4.1 M and +\$10.0 M and retail is the near-exact mirror (−\$3.6 M to −\$9.3 M); retail is on the losing side of 64–66% of manipulated cycles and bears 62–74% of the manipulated-cycle loss pool at every threshold. The MM cohort is a small bystander in manipulated cycles: its net PnL holds near −\$0.6 M throughout and its loss share stays modest (8–16%). Its behavioral profile — passive, delta-neutral quoting — is orthogonal to the manipulator’s directional-PnL threshold, so moving the threshold barely reclassifies it. In normal cycles the manipulator cohort’s PnL is near zero at moderate cells and only mildly negative at the tightest activity cuts ($n_{\text{cycle},\min} \geq 10$), confirming that manipulators do not systematically profit in normal cycles. The core finding — a large, retail-funded manipulated-cycle profit for the manipulator cohort — is robust to both threshold choices.

Table A5. Trader-type attribution sensitivity to manipulator thresholds. This table re-derives the full trader-type partition at each cell of a 4×5 threshold grid: the manipulator cohort is the set of wallets with $n_{\text{cycle},p} \geq n_{\text{cyc},\min}$ and net manipulated-cycle PnL $\geq \$_{\min}$; the MM cohort is the behavioral residual passing the same passive-rate, fills-per-active-cycle, and inventory criteria as the main text; retail/other is the final residual. Each panel reports one trader-type metric of Table 10 for all three cohorts, split across manipulated and normal cycles: *Panel A*, aggregate net PnL (\$M), plus the manipulator cohort size (Wallets); *Panel B*, the fraction of cycles in which the cohort is net negative; *Panel C*, the cohort's share of the gross loss pool, aggregated across cycles with pool-size weights. Within each regime the three cohort PnLs sum to zero and the three loss shares to 100%.

<i>Panel A. Net PnL (\$ M)</i>								
$n_{\text{cyc},\min}$	$_{\min}$	Wallets	Manipulated cycles (\$M)			Normal cycles (\$M)		
			Manip.	MM	Retail	Manip.	MM	Retail
3	\$500	2,276	+9.98	-0.65	-9.33	-0.92	+3.01	-2.08
3	\$1,000	1,456	+9.41	-0.64	-8.76	-0.19	+3.02	-2.83
3	\$2,000	857	+8.56	-0.62	-7.94	+0.28	+3.11	-3.39
3	\$5,000	414	+7.27	-0.59	-6.68	+0.68	+3.21	-3.90
3	\$10,000	221	+5.89	-0.55	-5.34	+0.31	+3.40	-3.70
5	\$500	2,158	+9.56	-0.65	-8.91	-0.94	+3.01	-2.07
5	\$1,000	1,389	+9.03	-0.64	-8.39	-0.30	+3.02	-2.72
5	\$2,000	821	+8.22	-0.62	-7.61	+0.09	+3.11	-3.20
5	\$5,000	397	+6.99	-0.59	-6.40	+0.52	+3.21	-3.73
5	\$10,000	210	+5.65	-0.55	-5.10	+0.16	+3.40	-3.56
10	\$500	1,949	+8.63	-0.65	-7.98	-1.43	+3.01	-1.58
10	\$1,000	1,275	+8.16	-0.64	-7.52	-0.89	+3.02	-2.13
10	\$2,000	751	+7.41	-0.62	-6.79	-0.53	+3.11	-2.58
10	\$5,000	356	+6.27	-0.59	-5.68	-0.04	+3.21	-3.18
10	\$10,000	181	+5.01	-0.55	-4.46	-0.31	+3.40	-3.09
20	\$500	1,673	+7.37	-0.65	-6.72	-1.43	+3.01	-1.58
20	\$1,000	1,130	+6.99	-0.64	-6.35	-1.05	+3.02	-1.97
20	\$2,000	672	+6.33	-0.62	-5.71	-0.75	+3.11	-2.36
20	\$5,000	315	+5.30	-0.59	-4.71	-0.33	+3.21	-2.88
20	\$10,000	151	+4.12	-0.55	-3.57	-0.54	+3.40	-2.86

Table A5. Trader-type attribution sensitivity to manipulator thresholds (continued).

<i>Panel B. % losing cycles</i>							
$n_{\text{cyc,min}}$	$\$_{\text{min}}$	Manipulated cycles			Normal cycles		
		Manip.	MM	Retail	Manip.	MM	Retail
3	\$500	35.8%	58.8%	66.3%	62.2%	37.6%	43.8%
3	\$1,000	36.8%	58.8%	66.0%	61.8%	37.7%	45.4%
3	\$2,000	38.1%	58.6%	65.4%	62.3%	37.7%	47.7%
3	\$5,000	38.2%	58.7%	65.7%	60.5%	37.5%	53.4%
3	\$10,000	40.2%	57.6%	65.6%	66.6%	37.2%	55.7%
5	\$500	36.0%	58.8%	65.8%	62.3%	37.6%	43.7%
5	\$1,000	36.9%	58.8%	65.7%	61.9%	37.7%	45.4%
5	\$2,000	38.0%	58.6%	65.3%	62.4%	37.7%	47.7%
5	\$5,000	38.2%	58.7%	65.5%	60.6%	37.5%	53.4%
5	\$10,000	40.3%	57.6%	65.2%	66.6%	37.2%	55.7%
10	\$500	36.7%	58.8%	65.0%	62.3%	37.6%	43.8%
10	\$1,000	37.2%	58.8%	65.1%	62.0%	37.7%	45.4%
10	\$2,000	38.4%	58.6%	64.8%	62.6%	37.7%	47.9%
10	\$5,000	38.6%	58.7%	65.2%	60.7%	37.5%	53.3%
10	\$10,000	40.7%	57.6%	64.8%	66.9%	37.2%	55.5%
20	\$500	37.2%	58.8%	64.4%	62.3%	37.6%	44.7%
20	\$1,000	37.3%	58.8%	64.4%	61.9%	37.7%	46.3%
20	\$2,000	38.7%	58.6%	64.4%	62.5%	37.7%	48.4%
20	\$5,000	38.5%	58.7%	65.0%	60.9%	37.5%	53.5%
20	\$10,000	40.7%	57.6%	64.3%	67.0%	37.2%	55.6%

<i>Panel C. Loss share</i>							
$n_{\text{cyc,min}}$	$\$_{\text{min}}$	Manipulated cycles			Normal cycles		
		Manip.	MM	Retail	Manip.	MM	Retail
3	\$500	17.8%	8.4%	73.8%	46.3%	9.1%	44.6%
3	\$1,000	18.7%	8.6%	72.6%	44.9%	9.6%	45.5%
3	\$2,000	19.7%	9.1%	71.2%	43.1%	10.6%	46.3%
3	\$5,000	20.3%	10.3%	69.4%	40.0%	12.4%	47.5%
3	\$10,000	20.1%	12.5%	67.3%	37.0%	15.3%	47.7%
5	\$500	18.4%	8.6%	73.0%	46.2%	9.3%	44.5%
5	\$1,000	19.3%	8.8%	71.9%	44.8%	9.8%	45.4%
5	\$2,000	20.2%	9.3%	70.5%	43.2%	10.8%	46.0%
5	\$5,000	20.7%	10.6%	68.7%	40.1%	12.7%	47.3%
5	\$10,000	20.6%	12.8%	66.6%	37.1%	15.5%	47.4%
10	\$500	19.1%	9.2%	71.6%	46.5%	9.9%	43.6%
10	\$1,000	19.9%	9.5%	70.5%	45.3%	10.4%	44.3%
10	\$2,000	20.8%	10.1%	69.1%	43.6%	11.6%	44.8%
10	\$5,000	21.5%	11.4%	67.1%	40.4%	13.5%	46.1%
10	\$10,000	21.5%	13.9%	64.7%	37.3%	16.3%	46.3%
20	\$500	19.8%	10.5%	69.7%	45.9%	10.7%	43.4%
20	\$1,000	20.5%	10.8%	68.7%	44.9%	11.2%	43.9%
20	\$2,000	21.3%	11.5%	67.2%	43.3%	12.4%	44.3%
20	\$5,000	22.2%	13.0%	64.8%	40.2%	14.3%	45.4%
20	\$10,000	22.2%	16.1%	61.7%	37.2%	17.1%	45.7%

F Push-Down Cohort Flow

The cohort-flow figure in Section 6.4.6 (Figure 12) conditions on *push-up* cycles, manipulated cycles whose final-10s spot push is upward, so that the directional accumulation reads as a positive ramp. This appendix repeats the same construction on the *push-down* mirror (manipulated cycles whose push is downward) in Figure A3.

The pattern is the sign-flipped mirror with the identical timing signature. The manipulator cohort ramps smoothly into a net *short* Up position (buying the Down token), reaching $\sim -9,500$ shares by close; MMs and retail take the long-Up side; and the Binance order flow is flat through the body and then steps down by $\sim \$1.75$ M in the final 10 s (cumulative $\sim -\$2$ M by close), the down-side counterpart of the $+\$1.7$ M up-push. The mirror is not exact in magnitude: both the final-10s Binance step ($\sim -\$1.75$ M vs $\sim +\$1.7$ M) and the manipulator cohort's cumulative directional position ($\sim -9,500$ vs $\sim +8,800$ shares) are marginally larger on the down side. But the linear-ramp-then-late-step timing is the same in both directions, confirming that the timing argument against the gamma-hedging channel (Section 6.4.6) is not an artifact of conditioning on the up direction.




Figure A3. Cumulative net position trajectory across the cycle, push-down vs. normal cycles, P3 BTC 5m. This figure plots the push-down counterpart of Figure 12; the solid line conditions on push-down cycles (labeled “manipulated cycles” in the legend), construction otherwise identical. Panels (a)–(c) plot the net Up-token holdings (shares) of the manipulator, MM, and retail/other cohorts; panel (d) plots the cumulative Binance spot order flow (\$). Dashed gray lines: normal cycles. Pink shading: the final-10 s region (seconds 290–300).

References

- Akey, Grégoire, Harvie, and Martineau (2026). *Who Wins and Who Loses in Prediction Markets? Evidence from Polymarket*. Working Paper 6443103. SSRN.
- Aoyagi, Jun and Yuki Ito (2025). “Coexisting Exchange Platforms: Limit Order Books and Automated Market Makers”. In: *Journal of Political Economy Microeconomics*.
- Arrow, Kenneth J. et al. (2008). “The Promise of Prediction Markets”. In: *Science* 320.5878, pp. 877–878.
- Aspembitova and Bentley (2022). “Oracles in Decentralized Finance: Attack Costs, Profits and Mitigation Measures”. In: *Entropy* 24.1, p. 60.
- Baltussen, Guido, Julian Terstegge, and Paul Whelan (2023). *The Derivative Payoff Bias*. Working Paper 4562800. SSRN.
- Barbon, Andrea and Angelo Ranaldo (2026). “On the Quality of Cryptocurrency Markets: Centralized versus Decentralized Exchanges”. In: *Management Science*.
- Bartlett, Robert P. and Maureen O’Hara (2026). *Adverse Selection in Prediction Markets: Evidence from Kalshi*. Working Paper 6615739. SSRN.
- Capponi, Agostino and Ruizhe Jia (2025). “Liquidity Provision on Blockchain-Based Decentralized Exchanges”. In: *Review of Financial Studies* 38.10, pp. 3040–3085.
- Capponi, Agostino, Ruizhe Jia, and Kanye Ye Wang (2025). “Maximal Extractable Value and Allocative Inefficiencies in Public Blockchains”. In: *Journal of Financial Economics* 172, p. 104132.
- Capponi, Agostino, Ruizhe Jia, and Shihao Yu (2026). “Price Discovery on Decentralized Exchanges”. In: *Review of Financial Studies*. Forthcoming; SSRN 4236993.
- Chiu, Jonathan and Thorsten V. Koepl (2019). “Blockchain-Based Settlement for Asset Trading”. In: *Review of Financial Studies* 32.5, pp. 1716–1753.
- Comerton-Forde, Carole and Tālis J. Putniņš (2011). “Measuring Closing Price Manipulation”. In: *Journal of Financial Intermediation* 20.2, pp. 135–158.

- Commodity Futures Trading Commission (Mar. 2026). *Staff Advisory on Event Contracts*. CFTC Staff Letter 26-08. Commodity Futures Trading Commission.
- Cong, Prasad, and Rabetti (2025). *Financial and Informational Integration Through Oracle Networks*. NBER Working Paper 33639. National Bureau of Economic Research.
- Daian et al. (2020). “Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges”. In: *IEEE Symposium on Security and Privacy (S&P)*.
- Della Vedova (2026). *Who Profits from Prediction Markets? Execution, not Information*. Working Paper 6191618. SSRN.
- Diercks, Katz, and Wright (2026). *Kalshi and the Rise of Macro Markets*. NBER Working Paper 34702. National Bureau of Economic Research.
- Dubach (2026). *The Anatomy of a Decentralized Prediction Market: Microstructure Evidence from the Polymarket Order Book*. Working Paper arXiv:2604.24366. arXiv.
- Duffie, Darrell and Piotr Dworczak (2021). “Robust Benchmark Design”. In: *Journal of Financial Economics* 142.2, pp. 775–802.
- Easley, O’Hara, Yang, and Zhang (2026). “Microstructure and Market Dynamics in Crypto Markets”. In: *Journal of Financial Markets*. SSRN 4814346.
- Glosten, Lawrence R. and Paul R. Milgrom (1985). “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders”. In: *Journal of Financial Economics* 14.1, pp. 71–100.
- Gómez-Cram, Guo, Jensen, and Kung (2026). *Prediction Market Accuracy: Crowd Wisdom or Informed Minority?* Working Paper 6617059. SSRN.
- Griffin, John M. and Amin Shams (2018). “Manipulation in the VIX?” In: *Review of Financial Studies* 31.4, pp. 1377–1417.
- Hanson, Robin and Ryan Oprea (2009). “A Manipulator Can Aid Prediction Market Accuracy”. In: *Economica* 76.302, pp. 304–314.

- Hanson, Robin, Ryan Oprea, and David Porter (2006). “Information Aggregation and Manipulation in an Experimental Market”. In: *Journal of Economic Behavior and Organization* 60.4, pp. 449–459.
- Hasbrouck, Joel, Thomas J. Rivera, and Fahad Saleh (2024). “An Economic Model of a Decentralized Exchange with Concentrated Liquidity”. In: *Management Science*.
- Hasbrouck, Joel, Fahad Saleh, and Thomas Rivera (2026). “The Need for Fees at a DEX: How Increases in Fees Can Increase DEX Trading Volume”. In: *Management Science*. Accepted.
- Klein, Frederik, Roman Kozhan, Ganesh Viswanath-Natraj, and Junxuan Wang (2023). *Price Discovery in Cryptocurrencies: Trades versus Liquidity Provision*. Working Paper. SSRN 4642411.
- Kumar, Praveen and Duane J. Seppi (1992). “Futures Manipulation with “Cash Settlement””. In: *Journal of Finance* 47.4, pp. 1485–1502.
- Lehar, Alfred and Christine A. Parlour (2025). “Decentralized Exchange: The Uniswap Automated Market Maker”. In: *Journal of Finance* 80.1, pp. 321–374.
- Lehar, Alfred, Christine A. Parlour, and Marius Zoican (2022). *Liquidity Fragmentation on Decentralized Exchanges*. Working Paper. SSRN 4267429.
- Makarov, Igor and Antoinette Schoar (2020). “Trading and Arbitrage in Cryptocurrency Markets”. In: *Journal of Financial Economics* 135.2, pp. 293–319.
- Mitts and Ofir (2026). *From Iran to Taylor Swift: Informed Trading in Prediction Markets*. Working Paper 6426778. SSRN.
- Ni, Sophie X., Neil D. Pearson, and Allen M. Potesman (2005). “Stock Price Clustering on Option Expiration Dates”. In: *Journal of Financial Economics* 78.1, pp. 49–87.
- Park, Andreas (2023). “The Conceptual Flaws of Decentralized Automated Market Making”. In: *Management Science* 69.11, pp. 6731–6751.
- Rhode, Paul W. and Koleman S. Strumpf (2008). *Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data*. NBER Working Paper 14377. National Bureau of Economic Research.

- Spatt, Chester (2014). “Security market manipulation”. In: *Annu. Rev. Financ. Econ.* 6.1, pp. 405–418.
- Tsang, Kwok Ping and Zichao Yang (2026). *The Anatomy of a Blockchain Prediction Market: Polymarket in the 2024 U.S. Presidential Election*. Working Paper arXiv:2603.03136. arXiv.
- Wolfers, Justin and Eric Zitzewitz (2004). “Prediction Markets”. In: *Journal of Economic Perspectives* 18.2, pp. 107–126.
- Zhang, Anthony Lee (2022). “Competition and Manipulation in Derivative Contract Markets”. In: *Journal of Financial Economics* 144.2, pp. 396–413.
- Zhu, Haoxiang (2014). “Do Dark Pools Harm Price Discovery?” In: *Review of Financial Studies* 27.3, pp. 747–789.