PRICE DISCOVERY ACROSS POLITICAL PREDICTION MARKETS: EVIDENCE FROM THE 2024 U.S. PRESIDENTIAL ELECTION*

EMREHAN AKTUĞ† *Sabancı University*

ORHAN TORUL‡ *Boğaziçi University*

June 2026

Abstract

During the 2024 U.S. presidential election, the same binary outcome—a Trump victory—was traded simultaneously on nine venues: a blockchain-based order book (Polymarket), two further exchanges (Betfair, Kalshi), and six dealer sportsbooks, across jurisdictions between which no mechanism enforces the law of one price. The result is a natural laboratory for asking where the efficient probability is produced, and whether that production matters beyond the betting ecosystem. Applying the Hasbrouck, Gonzalo–Granger, and Putniņš price-discovery decompositions to 5-minute data, we find a sharply hierarchical cross-section: the order-book exchanges Polymarket and Betfair jointly account for nearly 85% of information shares, no dealer sportsbook exceeds 11%, and the dealer books are consistently relegated to a following role. Benchmarking the prediction-market consensus against exchange-timestamped financial data, we show that it led S&P 500 futures throughout election night, led Trump Media's equity across the overnight gap, and priced the campaign's two weekend shocks while financial markets were closed. The newly CFTC-regulated entrant Kalshi attains nontrivial information shares within weeks, and leadership among the exchanges is horizon-dependent. We read the cross-section descriptively, since our design cannot fully separate venue organization from attention and liquidity. The cross-asset evidence does not share this limitation: prediction markets incorporate political information ahead of deep, professionally arbitraged financial markets, speaking directly to the regulatory debate over whether event-contract exchanges serve a legitimate price-discovery function.

Keywords: Price Discovery; Prediction Markets; Information Shares; Market Microstructure; Political Betting; Electoral Probabilities

JEL Classifications: D72; D83; G14; G18

*We are grateful to Koray Şimşek, Yiğit Atılgan, Burak Saltoğlu, Şerif Aziz Şimşir, and Hüseyin Gülen for their valuable feedback. All remaining errors are ours.

†Corresponding author. Address: Sabancı University Department of Economics, 34956, Tuzla, Istanbul, Türkiye.
E-Mail: emrehan.aktug@sabanciuniv.edu

‡Address: Boğaziçi University, Department of Economics, 34342, Bebek, Istanbul, Türkiye.
E-Mail: orhan.torul@bogazici.edu.tr

1 Introduction

During the 2024 U.S. presidential election, the probability of a Trump victory was bought and sold simultaneously on nine venues: a blockchain-based order book (Polymarket, $3.6 billion in election volume), two further order-driven exchanges (Betfair and Kalshi), and six dealer sportsbooks. These venues span continents, regulators, and trading technologies, yet each prices the same binary payoff, so the event effectively had as many prices as it had venues. Kalshi, the only CFTC-regulated venue, began listing election contracts in October 2024 after a protracted legal fight over whether political event contracts serve a legitimate economic purpose—a question now before regulators, and one to which our results speak directly.

What makes the setting distinctive is the severity of its frictions. In U.S. equity markets, fragmentation is disciplined by Regulation NMS, the National Best Bid and Offer, and high-frequency arbitrageurs who enforce the law of one price within milliseconds. None of these mechanisms exists here: U.S. residents cannot legally access Betfair; offshore and on-chain venues are barred to U.S. persons; capital is locked within each venue with no netting; and settlement ranges from minutes on-chain to days for sportsbook withdrawals. The same contract consequently trades at an average cross-venue spread of 6.2 percentage points. These are the limits to arbitrage of Shleifer and Vishny (1997) in unusually pure form, and they make the question that organizes this paper both well-posed and unresolved: with nothing tying the venues' prices together, *where is the efficient probability produced, and does its production matter beyond the betting ecosystem?*

We answer with the standard price-discovery toolkit of the equity microstructure literature—the Hasbrouck (1995) information share, the Gonzalo and Granger (1995) component share, and the Putniņš (2013) information leadership share—applied to 5-minute implied probabilities from all nine venues over January–November 2024. To our knowledge, this is the first application of the Hasbrouck decomposition to a multi-venue cross-section of political prediction markets. The framework is essential to the question: unlike a cross-correlation, which cannot distinguish a venue that *discovers* information from one that merely *copies* price changes quickly, the vector error-correction model identifies each venue's contribution to the common efficient probability.

Our four results are facets of this one question; the first three locate *where* discovery happens within the betting ecosystem, and the fourth asks whether it *matters* outside it. *First*, the order-book exchanges dominate: Polymarket and Betfair jointly account for nearly 85% of information shares in our clean specification, no dealer sportsbook exceeds 11%, and every other measure we compute places the sportsbooks in a following role. *Second*, leadership between the two exchanges is horizon-dependent—Polymarket leads at the 5-minute frequency, Betfair at the 1-hour frequency. We document this crossover but read it cautiously: both venues are order books, so it cannot reflect an architectural divide, and clientele, depth, and quote freshness are candidate explanations we cannot separate. *Third*, Kalshi attains nontrivial information shares within 32 days of its launch; the short, salience-laden window precludes strong inference, but rapid entry by a venue with a fraction of the incumbents' volume is itself notable. *Fourth*, and most consequentially, we step outside the betting ecosystem and benchmark the prediction-market consensus against exchange-timestamped financial data: it led S&P 500 futures throughout election night, led Trump Media's equity (DJT) across the overnight gap when the stock could not trade, and priced the campaign's two weekend shocks while financial markets were closed.

Because these financial markets are deep and professionally arbitraged, the lead cannot be dismissed as one venue merely trying harder than its betting rivals; it is direct, design-independent evidence that exchange-organized prediction markets perform genuine price discovery.

These findings contribute to two literatures. To the study of fragmentation and price discovery (Hasbrouck, 1995; O'Hara and Ye, 2011; Makarov and Schoar, 2020), we add an asset class in which cross-venue arbitrage is essentially absent and an external benchmark—liquid financial assets pricing the same event—that cross-venue studies rarely have. To the fast-growing literature on the 2024 prediction markets—closest is Ng et al. (2025) on Polymarket versus Kalshi—we add the betting exchange and dealer sportsbooks they omit, a horizon decomposition, and the cross-asset comparison; we develop this positioning in Section 2. Throughout, we interpret cautiously. The within-ecosystem cross-section is descriptive: because election contracts are a sideline for sportsbooks, their low information shares conflate venue organization with the attention and order flow a peripheral product attracts, and our design cannot fully separate the two. The cross-asset evidence carries no such caveat, and it is there that our strongest claim rests.

2 Institutional Background and Related Literature

2.1 Institutional Setting

Our nine venues fall into two organizational types: order-driven exchanges and dealer sportsbooks. The three exchanges—Polymarket, Betfair, and Kalshi—match buyers and sellers in a visible book of competing limit orders; the six sportsbooks instead quote odds as dealers. We take each type in turn. Polymarket operates on the Polygon blockchain from a base outside the United States. Contrary to a common characterization, Polymarket is *not* an automated market maker (AMM) during our sample. It launched in 2020 on a constant-function AMM (an LMSR-style design), but deprecated that mechanism and migrated to a hybrid-decentralized central limit order book (CLOB) over 2022–2023, well before our January 2024 sample begins.[1] Betfair Exchange, regulated by the UK Gambling Commission, and Kalshi, a CFTC-regulated designated contract market, are likewise centralized limit order books in which participants post and revise limit orders. The three exchanges thus differ mainly in jurisdiction, settlement technology, fee schedule, and clientele, not in how their prices form. Kalshi's entry on October 4, 2024, following the District Court ruling in *Kalshi v. CFTC*, is the key institutional event in our sample: it introduced a U.S.-regulated exchange into an ecosystem previously accessible only through offshore platforms.

Six dealer sportsbooks complete the sample: Pinnacle, Bovada, and BetOnline (sharp books with tighter margins) and William Hill, Unibet, and Everygame (retail-focused books with wider margins). Unlike the exchanges, sportsbooks quote odds set by internal risk models and revise them infrequently, generating the price staleness that is central to our empirical analysis. For all six, election betting is a peripheral product alongside

[1] Polymarket's order book is operated off-chain—limit orders are posted and matched by a central operator—while settlement and custody occur on-chain via smart contracts on Polygon. The migration from the AMM to the CLOB is documented in Ng et al. (2025). Throughout the 2024 election cycle, therefore, Polymarket prices are set by human limit orders, exactly as on a conventional exchange, rather than by an automated pricing function. In the data, Polymarket's recorded price changes in 1.1% of 5-minute bins, reflecting how often its marketable price moves outside of news events; this is below Betfair (5.4%) but an order of magnitude above the sportsbooks (0.1–0.4%).

a much larger sports business; we return to this point when interpreting their low information shares, because it means their measured performance reflects attention and effort as much as venue design. Table 1 in Section 3 provides a detailed comparison of all nine platforms.

2.2 Related Literature

Price discovery and market fragmentation. Hasbrouck (1995) introduced the information share (IS), measuring each market's contribution to the variance of the efficient price innovation from the moving-average representation of a VECM. Gonzalo and Granger (1995) proposed the component share (CS), an ordering-invariant alternative based on the error correction coefficients. Baillie et al. (2002) and De Jong (2002) compare the two, and Putniņš (2013) proposes the Information Leadership Share (ILS) to disentangle informational from speed advantages. These tools have been applied to equities (Hasbrouck, 1995), foreign exchange (Lehmann, 2002), futures (Baillie et al., 2002), ETFs, stock–option pairs (Chakravarty et al., 2004), and cryptocurrency exchanges (Makarov and Schoar, 2020), but not, to our knowledge, to prediction markets.

Our setting connects to the broader fragmentation literature. O'Hara and Ye (2011) show that fragmentation across U.S. equity venues improves market quality, while Foucault and Menkveld (2008) model how smart order routing affects price discovery across venues. Makarov and Schoar (2020) document large and persistent price differences across cryptocurrency exchanges, attributing them to capital controls and withdrawal frictions—an analogy that maps closely to our setting, where jurisdictional barriers and capital lockup similarly prevent cross-venue arbitrage.

Venue design: exchanges, dealers, and crypto venues. A long literature contrasts order-driven and quote-driven (dealer) markets and their price-discovery properties (Madhavan, 2000). Croxson and Reade (2014) study information incorporation in Betfair's betting exchange, documenting near-instantaneous adjustment to public news during football matches—the closest precedent for treating a betting exchange as a price-discovery venue. Because Polymarket settles on-chain, our setting also connects to work on decentralized and crypto venues, notably Lehar and Parlour (2021) on decentralized exchanges and automated market making; the persistent cross-venue crypto price gaps of Makarov and Schoar (2020), discussed above, are the closest existing analogue to the fragmentation we study. We stress, however, that Polymarket itself runs an order book rather than an AMM (Section 2); the relevant contrast in our data is between three order-book exchanges with different clienteles and a set of dealer sportsbooks, not between an AMM and a limit order book. In Appendix B we sketch a simple venue-choice framework that rationalizes why dealer books act as informational followers and why a newly opened regulated exchange can quickly attract informed order flow.

Prediction markets. A large literature examines prediction-market accuracy, calibration, and design (Wolfers and Zitzewitz, 2004; Manski, 2006; Snowberg et al., 2013; Page and Clemen, 2013; Arrow et al., 2008; Berg et al., 2008; Rothschild, 2009; Hanson et al., 2006; Lee et al., 2020). The consensus is that market-based forecasts are typically well calibrated and competitive with or superior to polls, though Manski (2006) cautions that equilibrium prices need not equal mean beliefs. Our question is orthogonal to this debate: rather than asking whether prediction-market prices are accurate, we ask which *type* of venue produces them.

The 2024 prediction-market literature. The 2024 election generated a wave of closely related work, against which we position ourselves explicitly. Ng et al. (2025) assemble transaction-level data on common contracts

across Polymarket, Kalshi, PredictIt, and Robinhood and estimate a 30-minute return VAR; they find that Polymarket leads Kalshi, that the lead concentrates in high-liquidity, high-volume periods, and that whale order imbalance predicts subsequent returns. Our paper is complementary: we add the betting exchange and dealer sportsbooks that they omit (a nine-venue cross-section), a horizon decomposition of leadership, the Kalshi-entry episode, and a cross-asset comparison with financial markets; they, in turn, observe the order flow and trade-level data that we do not. Other concurrent work studies the anatomy of Polymarket trading and betting-versus-polling accuracy (Yang and Tsang, 2026; Cutting et al., 2025), and Flynn and Tarkom (2025) show that betting markets incorporate political information faster than equity markets during 2024—a finding we revisit directly with high-frequency futures data in Section 5.7. Aktuğ and Torul (2026) and Snowberg et al. (2007) instead take prediction-market prices as inputs, using them to value partisan corporate alignment and to measure partisan economic effects. Relative to all of these, our distinct contribution is to bring the full Hasbrouck–Gonzalo-Granger apparatus to a multi-venue cross-section spanning three exchange types and six dealer books, and to validate the prediction-market consensus against an external financial benchmark.

2.3 Hypotheses

The organizational differences across our nine platforms motivate four predictions about the cross-section and dynamics of information shares, which we state ex ante from venue characteristics and, where possible, link to a simple venue-choice framework (Appendix B). We flag at the outset that the four are not on equal footing: H1 follows in part mechanically from the update-frequency gap documented below, and H3 partly describes the sample that motivated the study. We therefore treat the hypotheses as an organizing device for a descriptive cross-section rather than as sharp structural tests.

H1: Order-book exchanges lead dealer sportsbooks in price discovery. The three exchanges (Polymarket, Betfair, Kalshi) feature a visible book of competing limit orders, continuous trading, and participants who profit directly from correcting mispricing. Dealer sportsbooks instead quote odds from internal models, revise them infrequently, and treat election contracts as a sideline product, so they should track the efficient price set elsewhere with a lag (Proposition 3). Standard microstructure reasoning predicts that venues with continuous, transparent, incentive-aligned price formation lead price discovery (Madhavan, 2000; Hasbrouck, 1995). We stress that this prediction bundles venue design with effort and order flow, which our design cannot separate.

H2: A regulated exchange can rapidly attain nontrivial information shares. If exchange organization—rather than incumbency or accumulated liquidity—is what facilitates informed trading, then a newly launched, comparably-designed exchange (Kalshi) should attain nontrivial information shares within a short period. The framework shows that a regulated exchange that unlocks a previously excluded pool of informed traders can gain a meaningful informational presence even without matching the incumbent's volume (Proposition 2).

H3: Information leadership shifts toward the exchanges as election salience rises. If information shares reflect where informed traders concentrate, the rise in salience and volume over the campaign should reallocate price discovery from venues that led during the quiet early period toward the exchanges that attract sophisticated, high-frequency participants. We accordingly look for a regime shift around the July 2024 shocks. Because the timing of those shocks motivated the sample in the first place, we read this as a description of the episode rather than an out-of-sample prediction.

H4: Price-discovery leadership varies with the horizon over which it is measured. The three exchanges differ in immediacy, depth, and clientele: Polymarket is a crypto-native, 24/7 venue whose marketable price can move on a single aggressive order, whereas Betfair is a deep, long-established book used by experienced, higher-stakes participants. Such differences can produce a leader that depends on the sampling horizon—a faster-reacting venue contributing more at high frequency, a deeper venue's cumulative informed flow contributing more at lower frequency. We treat H4 as an *exploratory* empirical question rather than a structural prediction, precisely because both venues are order books and the candidate channels (speed, depth, clientele, quote freshness) cannot be separated with our data.

3 Data

3.1 Sources and Collection

Table 1 summarizes the nine platforms. From each we obtain 5-minute snapshots of the implied probability of a Trump victory over its coverage period, drawn from three types of source that we describe below. How a venue's price is sourced and sampled governs how often it appears to change, and that cross-section of update frequencies—as Section 3.3 shows—is the dominant force in the results.

Table 1: Data Summary

Platform	Architecture	Obs	Period	Coverage (%)	Prob Range
Polymarket	Order-book DEX (CLOB)	188,406	Jan 2023–Nov 2024	100.0	0.44–0.75
Betfair	Betting Exchange (P2P)	89,080	Dec 2023–Nov 2024	99.8	0.38–0.71
Pinnacle	Sharp Sportsbook	35,459	Jul 2024–Nov 2024	39.8	0.47–0.76
Bovada	Sharp Sportsbook	88,866	Dec 2023–Nov 2024	99.7	0.42–0.79
BetOnline	Sharp Sportsbook	71,195	Dec 2023–Nov 2024	79.9	0.45–0.75
William Hill	Retail Sportsbook	53,993	Dec 2023–Nov 2024	60.8	0.45–0.75
Unibet	Retail Sportsbook	85,785	Dec 2023–Nov 2024	96.2	0.44–0.80
Everygame	Retail Sportsbook	47,519	Dec 2023–Nov 2024	53.5	0.27–0.67
Kalshi	Regulated Exchange (CFTC)	9,358	Oct 4–Nov 5, 2024	—	0.47–0.66

Notes: Coverage indicates the percentage of 5-minute grid periods with valid observations during the primary sample period. Sportsbook decimal odds are converted to implied probabilities via $p = 1/\text{odds}$. Polymarket and Kalshi report implied probabilities directly. Pinnacle and Kalshi data reflect their respective mid-sample entries.

Sportsbooks and Betfair. For all six sportsbooks and for the Betfair exchange, we obtain decimal odds on the Trump-victory contract from *The Odds API* (`the-odds-api.com`), a commercial service that aggregates a panel of books and exchanges and timestamps each quote. We sample every five minutes, so each observation is the venue's most recent *posted* price as of the snapshot rather than a continuous quote stream: the aggregator refreshes faster than we sample, but the precise instant at which a book revises its odds is unobserved. Our data therefore record *which* venue carries a price and how often it changes, but not exactly when it moves—a limitation that, for venues revising only rarely, produces the stale pricing we analyze in Section 3.3. We convert

decimal odds to implied probabilities as $p = 1/d$, which does not remove the bookmaker's overround.[2]

Polymarket and Kalshi. Polymarket prices are taken from its on-chain order book on the Polygon blockchain (publicly recorded): each 5-minute observation is the prevailing marketable price of the Yes (Trump-win) outcome token, which pays \$1 on a Trump victory and is therefore read directly as a probability in $(0,1)$; short gaps in the on-chain record are handled by the forward-fill described in Section 3.2. Kalshi prices are last-transaction prices on its CFTC-regulated limit order book, obtained from `kalshi.com`. Unlike the aggregated odds, these two series come directly from the venue's own record of executed and resting limit prices, and they update far more often than the sportsbooks.

Financial benchmarks (DataBento). For the cross-asset analysis of Section 5.7, we obtain two exchange-timestamped financial series from DataBento: front-month CME E-mini S&P 500 futures (continuous `ES` contract, CME Globex MDP 3.0 feed) and the equity of Trump Media & Technology Group (`DJT`, Nasdaq TotalView-ITCH), both as one-minute open–high–low–close bars. These series carry venue-grade timestamps and serve as a high-quality external reference against which we benchmark the prediction-market consensus.

3.2 Sample Construction and Panels

We align the series to a 5-minute UTC grid by last-observation resampling and forward-fill gaps of up to 30 minutes (six bins). We define four analysis panels over the full January–November 2024 period: (i) a *clean panel* of six markets (excluding Pinnacle, which enters in July, and Everygame, which exhibits a staleness artifact documented in Section 5.2); (ii) a *common panel* of seven markets spanning the full sample (excluding Pinnacle); (iii) an *eight-market panel* including Pinnacle from July onward; and (iv) a *Kalshi-window panel* of nine markets during October 4–November 5, into which Kalshi enters. The clean panel is our primary specification; the others are used for robustness. The resampling and forward-fill choices interact with the very different update frequencies across venues in ways that shape the entire cross-section of information shares. Because this interaction is the single most important feature of the data, we devote the next subsection to it.

3.3 Update Frequency, Asynchronous Trading, and Forward-Fill

The nine venues update at rates that differ by more than an order of magnitude. The marketable price changes in 5.4% of 5-minute bins on Betfair and 1.1% on Polymarket, but in only 0.1–0.4% of bins on the sportsbooks (Table 2). This heterogeneity is not incidental; it is the dominant mechanical force shaping the cross-section of information shares, and it interacts with our data construction in two ways.

First, because the aggregated odds record *posted* prices, an unchanged price can mean either that the venue actively re-affirmed its quote or that it simply has not revised a stale one; the two are indistinguishable in the data. Second, the forward-fill procedure carries the last price forward whenever a venue does not post a new one within a 5-minute bin, producing a run of exactly-zero innovations in the VECM residual. When the venue finally updates, the accumulated catch-up appears as a single large jump whose squared magni-

[2] Because we study *changes* in implied probabilities rather than levels, a constant proportional margin differences out. Margins can nonetheless vary over time and across books—for instance, books may widen around major news—so the conversion is imperfect. We verify in Section 6 that the venue ranking is unchanged when we use raw rather than logit-transformed probabilities, and that the sportsbooks' subordinate role survives every specification we try.

tude can exceed the sum of the many small squared innovations that continuous updating would have generated. Mechanically, this redistributes estimated information shares: it depresses the measured contribution of frequently-updating venues and can *inflate* that of the stalest ones. The clearest casualty is Everygame, which updates in only 0.1% of bins and acquires an artifactual 21% information share (Section 5.2); it is excluded from the clean panel for precisely this reason.

We do not treat this as a nuisance to be footnoted but as a first-order threat to identification, and we confront it with four complementary devices: (i) excluding the stalest venue (the clean panel); (ii) a fresh-quote filter that conditions on active updating (Section 6.3); (iii) a Monte Carlo calibration of the staleness bias under the empirical update rates (Section 6.3); and (iv) re-estimation at coarser sampling frequencies, at which more venues have had time to update (Section 5.5). We report results only for patterns that survive all four. The reader should nonetheless keep this mechanism in mind: it means our information shares are best read as a *ranking* that is robust across these checks, not as precisely estimated quantities.

3.4 Logit Transformation

Implied probabilities are bounded on $[0,1]$ and exhibit trend-like dynamics over the campaign period. To ensure the innovation terms are not bounded and to improve the defensibility of the unit root assumptions, we apply the logit transformation $\ell_{i,t} = \ln(p_{i,t}/(1-p_{i,t}))$. All subsequent econometric analysis is conducted on the logit-transformed series. Section 6 demonstrates that our primary results are robust to the use of raw probabilities.

3.5 Descriptive Statistics

Table 2 reports summary statistics for the 5-minute first differences of the probabilities. The statistics are computed over the 10,491 complete-case observations in the eight-market panel (July 2024 onward, when Pinnacle enters); Pinnacle is omitted from the table for brevity. The VECM estimation in Section 5 uses the seven-market common panel (16,460 observations), which excludes Pinnacle to maintain a balanced sample from January onward.

Table 2: Descriptive Statistics of 5-Minute Probability Changes

Market	Mean (bp)	Std (bp)	Skewness	Kurtosis	Updates (%)	Roll Spread (bp)
Polymarket	−0.07	27.0	−2.25	2500	1.1	11.5
Betfair	−0.06	13.6	−0.57	25.4	5.4	7.3
Bovada	0.03	9.1	0.54	430	0.2	—
BetOnline	−0.09	11.8	−3.48	204	0.4	—
William Hill	−0.03	11.8	4.78	465	0.2	—
Unibet	−0.06	11.9	−2.08	359	0.2	—
Everygame	−0.04	51.5	−0.52	4712	0.1	33.2

Notes: Changes are 5-minute first differences of raw (untransformed) implied probabilities, in basis points. The VECM estimation in Section 5 uses logit-transformed series, which alter the variance structure. Updates (%) represents the fraction of periods in which the price changed. Roll Spread is the Hasbrouck-style effective spread. Sample: 10,491 complete-case observations from the eight-market panel (July 2024 onward); Pinnacle omitted from table for brevity.

Two features of the table bear on the analysis. First, the "Updates" column makes the heterogeneity of Section 3.3 concrete: sportsbooks change price in only 0.1–0.4% of bins, so most of their 5-minute "innovations" are mechanically zero. Second, Everygame's standard deviation (51.5 bp) is nearly twice the next-highest venue's (Polymarket, 27.0 bp), reflecting infrequent but large discrete adjustments; its residuals alone exhibit serial correlation (Ljung-Box $p < 0.01$), foreshadowing the staleness artifact of Section 5.2. The very high excess kurtosis of every series reflects the concentration of price changes in short bursts around news—the empirical counterpart of the jumpy, asynchronous updating that motivates the robustness battery in Sections 5.5 and 6.3.

3.6 Data Limitations

Three data limitations should be noted. First, we lack platform-level trading volume. While Polymarket's aggregate volume ($3.6 billion) is publicly available, comparable contract-level figures for the other venues are proprietary or unreported. Our interpretations of the mechanisms behind information shares—particularly the role of liquidity in attracting informed traders—therefore rest on institutional knowledge rather than direct measurement. The cross-asset benchmark of Section 5.7, which requires no venue volume, partially offsets this by validating the consensus against external financial data, and Ng et al. (2025) supply complementary transaction-level evidence.

Second, the nine series are not identical price objects. Polymarket and Kalshi prices are executed and resting limit prices, Betfair is an exchange price, and the sportsbook series are posted dealer odds embedding a time-varying overround. Combined with venue-specific frictions—jurisdictional access, settlement speed, withdrawal costs—these differences can sustain persistent cross-venue wedges, so the single-common-trend assumption underlying the VECM is a stronger one here than for a single security traded on linked exchanges. We confront this when choosing the cointegrating rank (Section 5.2) and report that our ranking survives imposing one fewer common relation.

Third, Kalshi's probability range during its 32-day window (0.47–0.66) is considerably narrower than the other venues'. Near $p = 0.5$ the logit transformation is approximately linear; at the extremes other platforms reach ($p \approx 0.38$ or 0.80) it amplifies variation substantially. This difference in dynamic range should be borne in mind when comparing Kalshi-window IS estimates to full-sample estimates.

3.7 Key Political Events

Five events anchor our event-window analysis (Figure 1); throughout, we report the Trump-victory probability, so the Democratic probability is its complement. The Biden–Trump debate (June 27) was the campaign's pivot: a halting performance began the slide in Biden's standing—Trump's victory probability edged up from about 60%—that would culminate in his withdrawal. The Trump assassination attempt (July 13), a Saturday-evening shock, raised Trump's probability by about 10 points within 48 hours. The Biden withdrawal and Harris endorsement (July 21) reversed part of that move as Harris consolidated Democratic support. Election Day (November 5) resolved the contract as the first returns arrived. We also include the Kalshi market launch (October 4), a market-structure event that reshaped the venue landscape. Figure 1 plots the probabilities

across the full sample with vertical lines at each event.

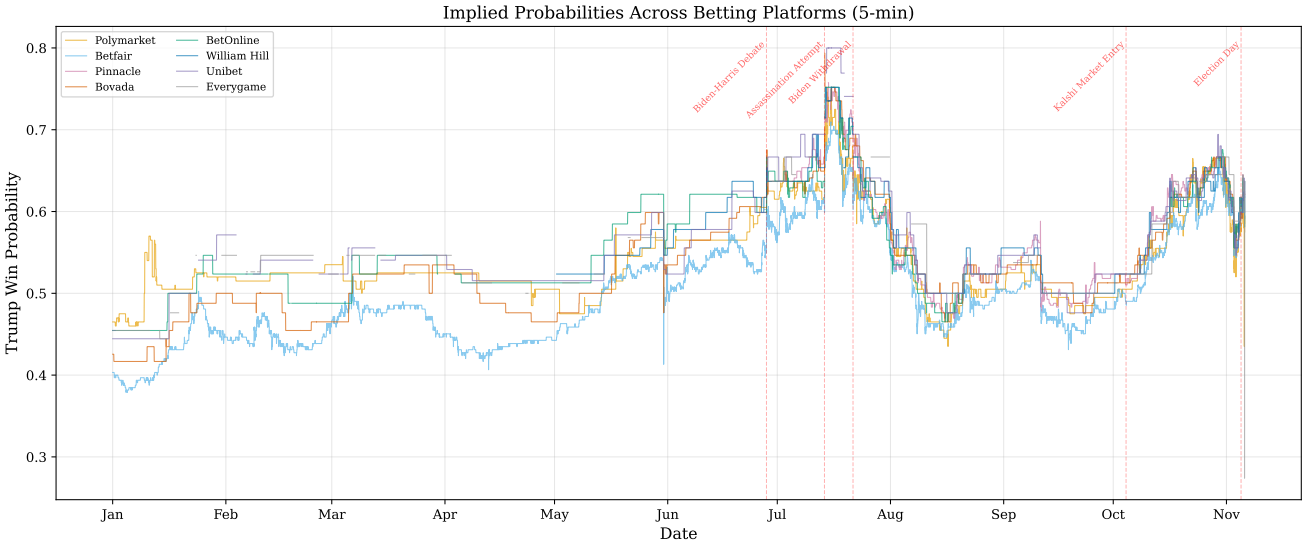


Figure 1: Implied Trump victory probabilities across eight platforms at 5-minute frequency, January–November 2024 (Kalshi excluded due to its shorter sample; Pinnacle enters in July). Vertical lines indicate key political events.

4 Econometric Methodology

Our empirical strategy uses one small core and a sequence of checks; we map them here so that a long list of tools is not mistaken for a long list of findings. The *core* is three standard measures—the Hasbrouck (1995) information share (IS), the Gonzalo and Granger (1995) component share (CS), and the Putniņš (2013) information leadership share (ILS)—which jointly characterize the price-discovery cross-section and separate informational from speed advantages. Every other tool answers one of two questions about that cross-section. *Is the ranking real, or an artifact?* Here we use weak-exogeneity and Granger tests (confirming the dealer books' follower role), the Hansen et al. (2011) Model Confidence Set (asking whether a single leader is statistically identified at all), and bootstrap, placebo, and Monte Carlo checks (quantifying sampling and staleness uncertainty). *Does the ranking depend on the horizon, or on the date?* Here we use the Diebold and Yılmaz (2014) and Baruník and Křehlík (2018) connectedness decompositions, re-estimation at coarser frequencies, rolling-window and DCC-GARCH estimates, and event-window and jump-timing analyses. A reader after only the headline can read the IS/CS/ILS cross-section (Section 5.2) and the cross-asset benchmark (Section 5.7); the rest is robustness and mechanism, not independent contributions.

4.1 Unit Root and Cointegration Tests

We test each logit-transformed series for a unit root using the Augmented Dickey–Fuller (ADF; Dickey and Fuller, 1979) test with automatic lag selection and the KPSS test (which has stationarity as its null). We test for cointegration using the Johansen (1991) trace and maximum-eigenvalue tests. The logit transformation

is essential here: raw probabilities are bounded on $[0,1]$ and cannot be unit root processes in the strict sense, whereas logit-transformed series are unbounded and can exhibit the persistent, mean-reverting-only-through-error-correction dynamics that the VECM framework requires.

With n markets tracking the same underlying binary event, the theoretical prediction is $n-1$ cointegrating relations (one common stochastic trend: the efficient probability). A lower rank would indicate that some venues track systematically different information sets—which, given the jurisdictional and settlement frictions documented above, is a real possibility rather than a mere statistical artifact. We therefore treat the theoretical rank as a benchmark to be checked against the data, not imposed blindly, and we report (Section 5.2) that our venue ranking is unchanged when we impose one fewer common relation than theory suggests.

4.2 Vector Error Correction Model

Let $\boldsymbol{\ell}_t = (\ell_{1t}, \ldots, \ell_{nt})'$ denote the vector of logit-transformed probabilities. If the series are cointegrated with rank r, the VECM representation is:

$$\Delta \boldsymbol{\ell}_t = \boldsymbol{\alpha}\boldsymbol{\beta}' \boldsymbol{\ell}_{t-1} + \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta \boldsymbol{\ell}_{t-i} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim (0, \boldsymbol{\Sigma}) \tag{1}$$

where $\boldsymbol{\alpha}$ $(n \times r)$ contains the error correction (adjustment) coefficients, $\boldsymbol{\beta}$ $(n \times r)$ the cointegrating vectors, and $\boldsymbol{\Gamma}_i$ the short-run dynamics. The lag order k is selected by the Bayesian Information Criterion (BIC) applied to the unrestricted VAR on first differences. The BIC selects $k = 17$ for the common seven-market panel, corresponding to 85 minutes of dynamics. This relatively high lag order reflects the stale-price dynamics generated by infrequently updating sportsbooks: when several venues update only in 0.1–0.2% of 5-minute bins (Table 2), the autocorrelation structure of their residuals extends over many lags, requiring a rich lag structure to achieve white-noise innovations. Section 6 confirms that the IS ranking is stable across $k \in \{13, 15, 17, 19, 21\}$.

4.3 Hasbrouck Information Shares

The Granger Representation Theorem implies the moving-average form:

$$\Delta \boldsymbol{\ell}_t = \boldsymbol{\Psi}(1)\boldsymbol{\varepsilon}_t + (1-L)\boldsymbol{\Psi}^*(L)\boldsymbol{\varepsilon}_t \tag{2}$$

where $\boldsymbol{\Psi}(1) = \boldsymbol{\beta}_\perp \left(\boldsymbol{\alpha}_\perp' \boldsymbol{\Xi} \boldsymbol{\beta}_\perp\right)^{-1} \boldsymbol{\alpha}_\perp'$ is the long-run impact matrix, with $\boldsymbol{\Xi} = \boldsymbol{I} - \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i$. When $r = n-1$, all rows of $\boldsymbol{\Psi}(1)$ are proportional, and we define $\boldsymbol{\psi}$ as the (normalized) common row vector representing each market's weight in the efficient price innovation.

The Hasbrouck (1995) information share of market j is:

$$\text{IS}_j = \frac{(\boldsymbol{\psi}' \mathbf{F}_{\cdot j})^2}{\boldsymbol{\psi}' \boldsymbol{\Sigma} \boldsymbol{\psi}} \tag{3}$$

where $\mathbf{F}$ is the lower-triangular Cholesky factor of the residual covariance matrix, $\boldsymbol{\Sigma} = \mathbf{F}\mathbf{F}'$, and $\mathbf{F}_{\cdot j}$ denotes its jth column. Because the Cholesky factorization depends on the ordering of variables, we compute IS_j for all $n!$ orderings and report upper and lower bounds, as well as the midpoint (Hasbrouck, 1995; Baillie et al., 2002).

For the nine-market Kalshi-window system (9! = 362,880), we sample 10,000 random orderings to approximate the bounds.

4.4 Gonzalo–Granger Component Shares

The Gonzalo and Granger (1995) permanent–transitory decomposition defines the permanent (efficient price) component as $\boldsymbol{\alpha}_\perp' \boldsymbol{\ell}_t$. The component share of market j is:

$$\mathrm{CS}_j = \frac{(\boldsymbol{\alpha}_\perp)_j}{\sum_{i=1}^{n} (\boldsymbol{\alpha}_\perp)_i} \tag{4}$$

This measure is point-identified (no ordering dependence) but can be negative, indicating systematic overshooting. It uses only the error correction coefficients and is thus independent of the residual covariance structure (Gonzalo and Granger, 1995; De Jong, 2002; Putniņš, 2013). To disentangle the two channels through which a market can lead—receiving information first versus adjusting more rapidly—we also compute the Putniņš (2013) Information Leadership Share, $\mathrm{ILS}_j = \mathrm{IS}_j / \mathrm{CS}_j$. Markets with ILS > 1 have an *informational* advantage (they contribute more to innovation variance than their error-correction weight implies), while ILS < 1 indicates a *speed* advantage (rapid adjustment without proportional informational contribution). The ILS is defined only for markets with $\mathrm{CS}_j > 0.01$; markets with near-zero CS produce unstable ratios.

4.5 Modified Information Shares

The ordering dependence of Hasbrouck IS is a well-known limitation. Following Lien and Shrestha (2009), we also compute the *Modified Information Share* (MIS), which replaces the covariance matrix $\boldsymbol{\Sigma}$ with the correlation matrix $\mathbf{R}$ in the Cholesky decomposition:

$$\mathrm{MIS}_j = \frac{\left(\boldsymbol{\psi}' \mathbf{D}^{1/2} \mathbf{M}_{\cdot j}\right)^2}{\sum_{k=1}^{n} \left(\boldsymbol{\psi}' \mathbf{D}^{1/2} \mathbf{M}_{\cdot k}\right)^2} \tag{5}$$

where $\mathbf{D}^{1/2} = \mathrm{diag}(\sqrt{\sigma_{11}}, \ldots, \sqrt{\sigma_{nn}})$ and $\mathbf{M}$ is the lower-triangular Cholesky factor of the correlation matrix $\mathbf{R} = (\mathbf{D}^{1/2})^{-1} \boldsymbol{\Sigma} (\mathbf{D}^{1/2})^{-1}$. The MIS is approximately ordering-invariant when cross-market residual correlations are small (Lien and Shrestha, 2009); in our data, where contemporaneous correlations of 5-minute changes are near zero (0.002–0.051), the approximation is tight and MIS values are nearly identical across orderings. We use MIS as an alternative numerator in the Putniņš (2013) ILS computation.

4.6 Weak Exogeneity Tests

A market is weakly exogenous in the VECM if its row of $\boldsymbol{\alpha}$ is zero—meaning it does not adjust to cointegrating errors and is thus a pure price leader. For each market j, we test $H_0 \colon \alpha_{j,1} = \cdots = \alpha_{j,r} = 0$ using a Wald χ^2 test with r degrees of freedom. Rejection indicates the market participates in error correction; the magnitude of the test statistic reveals the *intensity* of adjustment, providing a formal complement to the IS and CS measures.

4.7 Connectedness Analysis

We complement the VECM-based price discovery measures with the Diebold and Yılmaz (2012, 2014) connectedness framework, which constructs a *directed information flow network* from generalized forecast error variance decompositions (GFEVD; Pesaran and Shin, 1998). Unlike the IS, which identifies each market's share of the efficient price innovation, the connectedness table shows *who transmits information to whom*. For each market j, we compute directional spillovers: FROM (share of j's forecast error variance due to other markets), TO (share of other markets' variance due to j), and NET = TO − FROM. A positive NET identifies net information transmitters. The total connectedness index (TCI) summarizes system-wide integration. We estimate the GFEVD at a 12-period (1-hour) forecast horizon using the VMA representation derived from the VECM.

Following Baruník and Křehlík (2018), we decompose the connectedness measures into frequency bands—short-run (<30 minutes), medium-run (30 minutes to 2 hours), and long-run (>2 hours)—using spectral methods applied to the VMA representation. This provides a within-model test of horizon-dependent leadership, without requiring re-estimation at different sampling frequencies.

4.8 Bootstrap Confidence Intervals

We construct confidence intervals using two complementary approaches. Our primary method is the stationary bootstrap of Politis and Romano (1994) with automatic block length selection following Politis and White (2004) and Patton et al. (2009). Unlike the fixed-block bootstrap, the stationary bootstrap draws geometrically distributed block lengths, yielding a strictly stationary resampling scheme. The data-driven optimal block length (22 periods for the common panel) adapts to the dependence structure of each series. We apply the bias-corrected accelerated (BCa) correction of Efron and Tibshirani (1993) to address the first-order bias arising from the bounded $[0,1]$ support of IS estimates.

As a robustness check for conditional heteroskedasticity (volatility spikes around political events), we also implement a recursive-design wild bootstrap (Cavaliere and Taylor, 2008; Gonçalves and Kilian, 2004). Each replication multiplies VECM residuals by Rademacher random variables (± 1), preserving the conditional variance profile, then recursively constructs bootstrap levels using the estimated VECM dynamics.

4.9 Time-Varying and Event-Window Analysis

To examine how information leadership evolves, we estimate rolling-window VECMs with a 2-week window (4,032 five-minute periods) stepped by 1 day (288 periods), yielding approximately 44 overlapping windows across the common sample. The 2-week window represents a trade-off: shorter windows would capture finer dynamics but yield imprecise estimates in a 7-market VECM with 17 lags (119 short-run parameters per equation, 833 system-wide); longer windows would smooth over the event-driven regime shifts that are central to our analysis. For event-window analysis, we estimate separate VECMs in windows of 7 or 14 days centered on each of the five key political events.

5 Results

5.1 Preliminary Tests

Unit roots. Table 3 reports unit root test results for the seven common-panel markets. At the 5% level, all logit-transformed series fail to reject the ADF null of a unit root (p-values 0.08–0.32) and reject the KPSS null of stationarity (p-values < 0.01). First differences uniformly reject the ADF null (p-values < 0.001) and fail to reject KPSS. We conclude that all series are integrated of order one, $I(1)$.

Table 3: Unit Root Tests on Logit-Transformed Probabilities

	Levels			First Differences		
Market	ADF Stat	ADF p	KPSS Stat	ADF Stat	ADF p	KPSS Stat
Polymarket	−1.922	0.322	5.203**	−20.261	0.000	0.152
Betfair	−2.446	0.129	6.707**	−35.274	0.000	0.150
Bovada	−2.544	0.105	5.208**	−93.890	0.000	0.160
BetOnline	−2.540	0.106	2.902**	−54.953	0.000	0.184
William Hill	−2.670	0.079	4.231**	−128.285	0.000	0.118
Unibet	−2.616	0.090	3.360**	−53.153	0.000	0.154
Everygame	−2.525	0.110	5.085**	−24.483	0.000	0.134

Notes: ADF tests include an intercept (no trend) with automatic lag selection (AIC). KPSS tests use the Newey–West automatic bandwidth with a level specification. ** indicates rejection at the 1% level. The KPSS critical value at 5% is 0.463. Sample: 16,460 complete-case observations at 5-minute frequency.

Cointegration. The Johansen trace test applied to the seven-market system (with BIC-selected lag order $k = 17$) rejects the null of rank $\leq r$ for $r = 0, 1, 2, 3, 4$ at the 5% level, consistent with at least five cointegrating relations. Given that all seven markets track the same binary event, we impose the theoretically motivated rank $r = n - 1 = 6$ (one common stochastic trend). This choice warrants discussion: the statistical evidence supports rank 5, not 6. The gap could reflect genuine informational segmentation—some markets may track partially distinct information sets—or finite-sample distortions of the trace test in a system with stale-price dynamics. We maintain $r = n - 1$ as our baseline because it corresponds to the strongest theoretical prior (a single underlying binary event), but Section 6 demonstrates that results with $r = 5$ preserve the qualitative ranking of venues.

For the nine-market Kalshi-window system ($k = 3$, BIC-selected), both trace and maximum-eigenvalue tests indicate rank = 8, exactly $n - 1$.

5.2 Information Shares and Leadership

Full sample.

Table 4 reports Hasbrouck information shares and Gonzalo–Granger component shares for the six-market clean panel (excluding Everygame), estimated on 16,460 complete-case observations.[3]

[3] Everygame, a retail sportsbook that updates in only 0.1% of 5-minute bins, generates an artifactual IS of 21% due to price staleness (see Section 6.3). Its residuals exhibit serial correlation (Ljung-Box $p < 0.01$), confirming model misspecification. The 7-market specification including Everygame is reported in Appendix Table A.1 for completeness; the exchange-dominance finding is unchanged.

Table 4: Information Shares and Component Shares: Full Sample, Clean Panel (January–November 2024)

	Hasbrouck IS				
Market	Lower	Mid	Upper	CS	Arch.
Polymarket	0.375	**0.475**	0.576	0.370	DEX
Betfair	0.267	0.371	0.474	0.414	Exchange
BetOnline	0.027	0.110	0.193	0.175	Sharp
Bovada	0.006	0.059	0.112	0.111	Sharp
Unibet	0.004	0.044	0.084	0.078	Retail
William Hill	0.000	0.024	0.048	−0.148	Retail

Notes: 6-market VECM excluding Everygame. Hasbrouck bounds computed over all 6! = 720 Cholesky orderings. CS (Gonzalo–Granger component shares) are point-identified. VECM estimated with $k = 17$ lags and cointegrating rank $r = 5$.

The cross-section is sharply hierarchical. The two order-book exchanges—Polymarket (IS midpoint 47.5%) and Betfair (37.1%)—together account for the vast majority of the common efficient-price innovation, while no dealer sportsbook exceeds 11% and three of the four sit below 6%. We read the table primarily as an exchange-versus-dealer split rather than as a precise ranking of individual books: the sportsbook point estimates are small and imprecisely estimated, and we lean on the architecture-grouped decomposition of Section 5.4—which collapses the books into a single bloc and is far more parsimonious—as the cleaner statement of the headline. This pattern is the cross-sectional content of H1, with the caveat of Section 2 firmly in view: because election contracts are a sideline for the sportsbooks, their low shares conflate venue design with limited attention and order flow. The component shares broadly confirm the exchange-over-dealer hierarchy, with a different internal ordering—Betfair (0.41) slightly exceeds Polymarket (0.37) in permanent-component weight, and William Hill receives a negative weight (−0.15), indicating systematic overshooting relative to the efficient price.

Bootstrap confidence intervals are wide. In the 7-market specification (Appendix Table A.1), Polymarket's 90% bootstrap CI spans [1.8%, 84.4%]. This width is typical for Hasbrouck IS estimation, as the Cholesky decomposition is sensitive to small perturbations in the residual covariance matrix $\boldsymbol{\Sigma}$ (Baillie et al., 2002). The bootstrap CIs should not be confused with the Cholesky ordering bounds, which are narrow for Polymarket [37.5%, 57.6%] and confirm that its point estimate is not an artifact of variable ordering. The IS *ranking* should be interpreted as a pattern supported by consistent point estimates across multiple robustness checks (Section 6), not as a set of precisely estimated quantities.

Figure 2 displays the Hasbrouck information shares with bounds.

The Kalshi window. Table 5 reports results for the nine-market system during October 4–November 5, 2024 (5,658 complete-case observations).

Polymarket (44%) and Kalshi (41%) jointly account for 85% of price discovery. The Hasbrouck bounds are tight (upper–lower spread of 3–7 percentage points for the leading markets), indicating that this ordering is robust to the Cholesky factorization. All traditional sportsbook IS midpoints fall below 5%. Betfair's IS falls sharply from 37.1% in the full-sample clean panel to 5.7% here. This decline reflects the joint effect of expanding from a 6-market to a 9-market system, a much shorter lag order ($k = 3$ vs. $k = 17$), and the concentration of

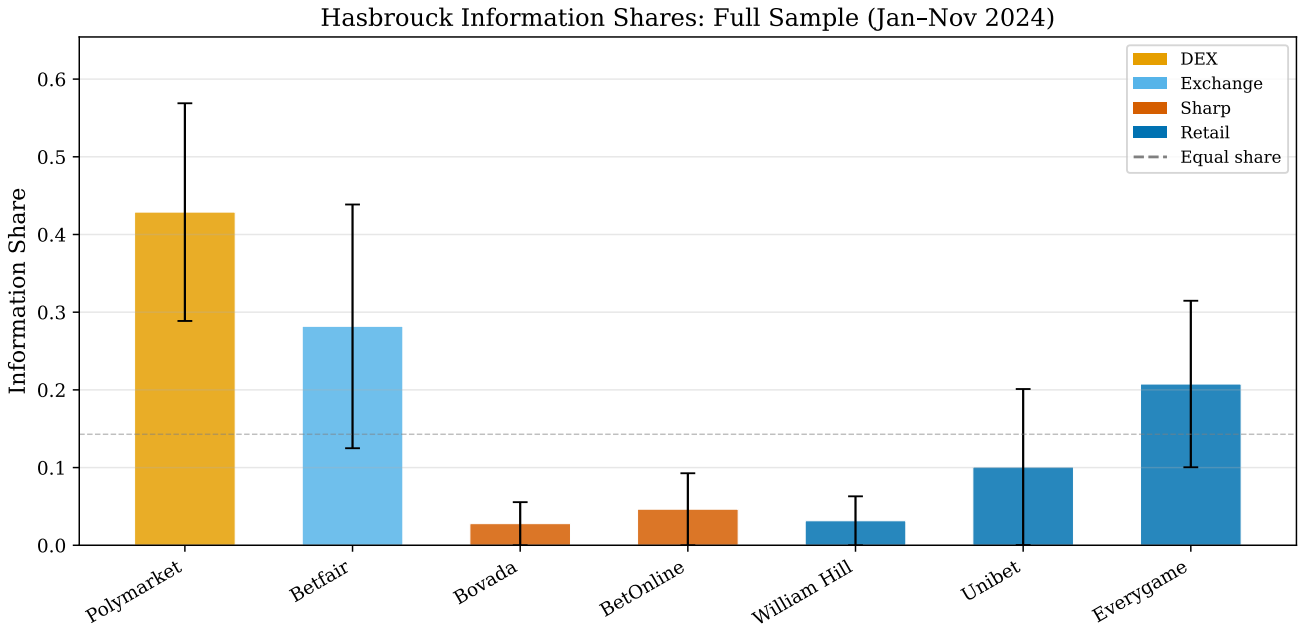


Figure 2: Hasbrouck information shares for the seven-market system (full sample, including Everygame for visual completeness; the primary clean-panel specification in Table 4 excludes Everygame). Error bars show upper and lower bounds across all Cholesky orderings. Markets are color-coded by architecture.

Table 5: Information Shares: Kalshi Window (October 4–November 5, 2024)

	Hasbrouck IS				
Market	Lower	Mid	Upper	CS	Arch.
Polymarket	0.403	**0.438**	0.474	0.298	DEX
Kalshi	0.378	**0.409**	0.441	0.242	Reg. Exch.
Betfair	0.035	0.057	0.079	0.205	Exchange
Pinnacle	0.025	0.031	0.036	0.224	Sharp
Bovada	0.008	0.018	0.027	0.133	Sharp
Everygame	0.044	0.048	0.053	−0.061	Retail
William Hill	0.001	0.002	0.003	0.032	Retail
BetOnline	0.000	0.002	0.003	−0.058	Sharp
Unibet	0.000	0.000	0.001	−0.016	Retail

Notes: Hasbrouck bounds approximated using 10,000 random orderings (of 9! = 362,880 total). CS sum to 1. VECM with $k = 3$ lags and rank $r = 8$.

information shares in Polymarket and Kalshi during the final month of the campaign.

The CS tell a complementary story: during the Kalshi window, the three exchange-based venues (Polymarket 0.298, Kalshi 0.242, Betfair 0.205) collectively receive 74% of the permanent-component weight, with Pinnacle (0.224) capturing a meaningful share that its low IS does not suggest. This divergence (Pinnacle having a high CS but low IS) indicates that Pinnacle's price is little adjusted toward the other venues—so it carries substantial permanent-component weight—yet its own innovations contribute relatively little variance. The low Putniņš ILS (= IS/CS ≈ 0.14) that follows implies a speed, rather than informational, contribution: Pinnacle's permanent-component weight derives from its slow error correction rather than from generating new information. This is informative for identification. Pinnacle is the sharpest book in the market—it accepts large stakes, moves its lines aggressively on information, and copies others least of all the venues here, so it is precisely the book *least* vulnerable to the charge that sportsbooks simply neglect election contracts. That even Pinnacle behaves as a pure follower (substantial permanent-component weight, negligible innovation share) suggests that the exchange-over-dealer gap reflects the price-formation mechanism—continuous, anonymous order matching versus periodically revised dealer quotes—and not merely effort or attention. We do not over-read a single venue, but Pinnacle is the hardest case for the "sideline-product" explanation, and it still follows.

That Kalshi attained nontrivial IS point estimates within weeks of launch, consistent with H2 and Proposition 2, carries suggestive regulatory implications, though the wide confidence intervals below caution against reading Kalshi's IS as on a par with Polymarket's. The conceptual framework identifies a possible mechanism: a regulated exchange that opens access to a previously excluded pool of informed traders can gain meaningful informational presence even without matching the incumbent's volume. Three important caveats apply: (i) this finding rests on a 32-day window during the campaign's most informationally intense period, when all markets should have been most efficient; (ii) block bootstrap confidence intervals (Appendix Table A.2) reveal wide 90% bounds for Polymarket [1.3%, 64.8%] and Kalshi [0.8%, 41.7%], reflecting severe parameter uncertainty over such a short estimation window; and (iii) the short sample window makes the estimates sensitive to idiosyncratic events during October 2024. The event-window analysis in Section 5.3 provides complementary evidence that Kalshi's information share built up gradually over its first month.

Information leadership. To better distinguish *informational* from *speed* advantages, Table 6 reports the Putniņš (2013) Information Leadership Share (ILS = IS/CS) for markets with positive CS, separately for the full sample and the Kalshi window.

The ILS reveals a sharp shift between the full sample and the Kalshi window. In the full sample, *all* markets with positive CS exhibit speed advantages (ILS < 1): Polymarket (0.40), Betfair (0.24), and BetOnline (0.16). Over the full 10 months, these markets' high CS values reflect rapid adjustment to equilibrium rather than the injection of new information. That even exchange-based venues exhibit speed rather than informational advantages over the full sample is consistent with the low information intensity during the early campaign months (January–June), when the election outcome was less salient and informed trading was sparse.

In the Kalshi window, Polymarket (ILS = 1.47) and Kalshi (ILS = 1.69) flip to *informational* advantages (ILS > 1), while all traditional venues remain below 0.30. During this final month of the campaign, when information intensity and trading volume were at their peak, the two exchange-based venues were not merely adjust-

Table 6: Information Leadership Share (ILS = IS/CS)

	Full Sample			Kalshi Window		
Market	IS	CS	ILS	IS	CS	ILS
Polymarket	0.429	1.061	0.40	0.438	0.298	**1.47**
Betfair	0.282	1.171	0.24	0.057	0.205	0.28
BetOnline	0.046	0.295	0.16	—	—	—
Bovada	0.028	0.006	—	0.018	0.133	0.13
Pinnacle	—	—	—	0.031	0.224	0.14
Kalshi	—	—	—	0.409	0.242	**1.69**

Notes: ILS = IS/CS; only defined for markets with CS > 0.01. ILS > 1 indicates an informational advantage (IS exceeds what error correction alone predicts); ILS < 1 indicates a speed advantage. Markets with negative or near-zero CS are omitted. CS values sum to 1 by construction; individual values can exceed 1 when other markets have large negative CS (Unibet −0.84, Everygame −0.69, William Hill −0.01 in the full sample), indicating those venues systematically overshoot relative to the efficient price. Full-sample column uses the 7-market VECM (including Everygame; see Appendix Table A.1). Kalshi window: 9-market VECM.

ing faster but were generating permanent price innovations that other venues subsequently tracked. Betfair, despite retaining a meaningful CS weight (0.205), exhibits only a speed advantage (ILS = 0.28), suggesting that it processes information originating at Polymarket and Kalshi rather than independently discovering it during this period.

Two implications follow. First, the exchange-versus-sportsbook hierarchy is not merely a mechanical artifact of price-updating frequency. If it were, the ILS would be uniformly below 1 across all venues and all windows, with faster venues leading only because they adjust sooner. The ILS exceeding 1 for Polymarket and Kalshi during the Kalshi window indicates that these venues attract informed order flow that generates permanent price innovations. Second, Kalshi achieved an ILS above one, comparable to Polymarket's, within weeks of launch (1.69 vs. 1.47 on the Hasbrouck measure; the two are essentially tied, 1.57 vs. 1.59, on the ordering-invariant MIS measure of Section 5.8). This is consistent with Proposition 2: the regulated exchange may have unlocked a pool of U.S.-based informed traders who were previously excluded from all other venues.

5.3 Time-Varying Leadership and Event Windows

Rolling information shares. Figure 3 displays rolling Hasbrouck information shares (using the default variable ordering from the data panel, 7-day smoothed) across the common sample. The single-ordering IS is used for computational tractability in the rolling analysis; the full-sample results in Table 4 confirm that midpoints across all orderings yield the same ranking. Consistent with H3, the dynamics reveal three distinct regimes:

1. **May–July (Traditional dominance):** Betfair and William Hill each hold 30–45% of price discovery. Polymarket contributes 5–10%. The traditional, established venues lead during the low-volatility early campaign period.
2. **July–August (Regime shift):** The assassination attempt and Biden withdrawal trigger a sharp reallocation. William Hill's share collapses from 40% to below 10%. The landscape becomes more fragmented.
3. **September–November (Prediction market dominance):** Polymarket's share rises steadily to 30–45%, es-

tablishing clear leadership through Election Day. This coincides with Polymarket's surge in trading volume and media coverage.

These regimes reconcile two facts that might otherwise look contradictory. Polymarket has the highest *full-sample* information share (47.5%, Table 4), yet it contributes only 5–10% in the rolling windows of the early campaign. The reconciliation is that the full-sample decomposition implicitly weights each period by its contribution to total innovation variance: the high-volatility final months—when Polymarket leads—carry the overwhelming majority of that variance, while the long, quiet early period contributes little. An equal-weighted average of the rolling windows instead gives the placid early months as much weight as the decisive final ones, and on that basis Betfair, which led early and at lower frequencies, comes out ahead (Section 5.8). The two summaries are not in conflict; they answer different questions—who led when it mattered most, versus who led on an average day—and we report both.

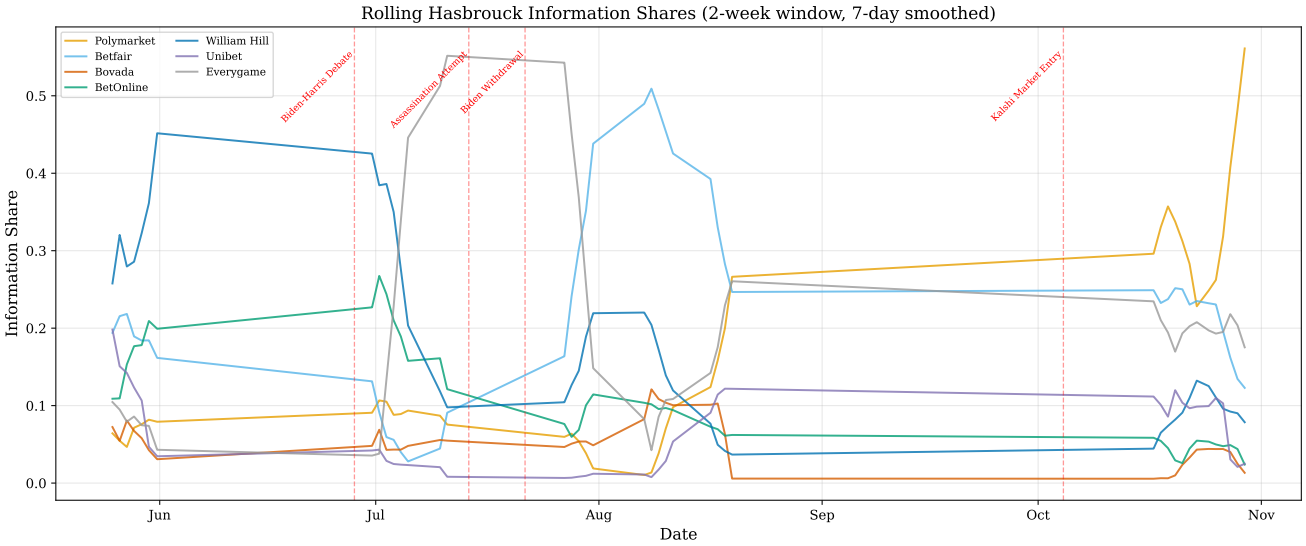


Figure 3: Rolling Hasbrouck information shares (2-week window, 1-day step, 7-day moving average). Vertical dashed lines indicate key political events. Note: rolling windows are backward-looking from each date.

Formal structural break tests confirm these visual impressions. Chow-type tests (Table 7) reject the null of parameter stability at all three event dates. The assassination attempt produces the largest F-statistic ($F = 6.59$), consistent with this exogenous weekend shock having the most disruptive effect on the cross-market price discovery structure.

Event windows. Table 8 reports Hasbrouck midpoints from VECM estimation in windows around each key event.

The VECM is estimated on short windows (7 or 14 days) with limited degrees of freedom, so the event-window IS should be read as descriptive rather than precisely estimated. The debate window (June 27) shows traditional sportsbooks leading: William Hill (32%) and Betfair (24%) dominate, while Polymarket contributes just 3%, consistent with the debate occurring during the period of traditional-platform dominance when Polymarket's volume was still building. The assassination attempt (July 13), an exogenous weekend shock, shows Polymarket surging to 24%. The shooting occurred on a Saturday evening; prediction markets (which trade

Table 7: Structural Break Tests

Event	n (markets)	F-stat	p-value	Verdict
Assassination Attempt (Jul 13)	6	6.590	< 0.001	Reject
Biden Withdrawal (Jul 21)	6	3.177	< 0.001	Reject
Kalshi Market Entry (Oct 4)	7	4.105	< 0.001	Reject

Notes: Chow-type F-tests comparing full-sample VECM residuals to separate pre/post-break VECM residuals. July events use a 6-market panel (excluding Everygame, which has only 4% coverage during July 11–22). The Kalshi entry test uses the full 7-market panel.

Table 8: Hasbrouck Information Share Midpoints by Event Window

Event	PM	BF	BOV	BO	WH	UNI	EG
Biden–Harris Debate	0.031	0.239	0.142	0.200	**0.315**	0.162	0.177
Assassination Attempt	0.237	0.030	0.129	0.085	0.029	0.060	0.513†
Biden Withdrawal	**0.371**	**0.406**	0.163	0.019	0.045	0.196	0.055
Kalshi Entry‡	0.175	0.114	**0.394**	0.002	0.029	0.006	0.091
Election Day	0.145	0.203	0.044	0.021	0.207	0.052	0.330†

Notes: PM = Polymarket, BF = Betfair, BOV = Bovada, BO = BetOnline, WH = William Hill, UNI = Unibet, EG = Everygame. Bold indicates the leading market(s) in each window. Window widths vary: 7 days (debate, election), 14 days (assassination, withdrawal, Kalshi entry). Everygame's high IS in these windows likely reflects the estimation artifact discussed in Section 5.2. Kalshi entry window uses the 9-market panel; Kalshi IS = 0.190, Pinnacle IS = 0.020 (not shown in columns).

24/7) adjusted immediately, while sportsbook lines remained stale until oddsmakers manually updated.[4] The Biden withdrawal (July 21) shows Polymarket (37%) and Betfair (41%) jointly accounting for 78% of price discovery during the campaign's most consequential single event.

The Kalshi Market Entry event window (centered on October 4) shows a gradual build-up. In the narrow window around its launch, Kalshi's IS is just 19%—lower than Bovada (39%) and comparable to Polymarket (18%). Yet the full 32-day Kalshi-window VECM (Table 5) shows Kalshi at 41%. This suggests that Kalshi's information share *built up* over the course of October, rather than appearing instantaneously. The gradual ramp-up is consistent with informed traders migrating to the new platform as it established liquidity and demonstrated reliable execution.

5.4 Architecture-Grouped IS Decomposition

To test the headline claim directly with a low-dimensional system, we collapse the seven markets into architecture groups: DEX (Polymarket), Exchange (Betfair), and Sportsbook (average of five sportsbooks), and estimate a 3-market VECM(5, $r = 2$) (Appendix Table A.9). With only three series, the system is parsimonious and the IS estimates are transparent. The DEX group captures 64.2% of information shares [52.2%, 76.2%], the Exchange group 33.0% [21.0%, 45.0%], and the Sportsbook composite just 5.1% [0.1%, 10.2%]. Gonzalo–Granger shares confirm the pattern: the DEX and Exchange groups are information leaders (66.2% and 58.1%), while

[4] Everygame's 51% IS in this window is the staleness artifact discussed in Section 5.2: it updated only twice during the window.

the Sportsbook composite is a systematic follower (−24.3%). Excluding Everygame from the Sportsbook composite raises it modestly to 13.3%, but the exchange dominance is unchanged (DEX 62.6%, Exchange 33.0%).

A finer 4-group decomposition, separating Sharp sportsbooks (Bovada, BetOnline) from Retail (William Hill, Unibet, Everygame), reveals that the DEX still dominates (52.5%), followed by the Exchange (22.4%), with the Sharp and Retail composites receiving 17.3% and 16.2% respectively. This confirms that even pooling sportsbooks by information reputation does not dislodge the exchange hierarchy. The architecture-grouped results demonstrate that the exchange-over-sportsbook finding is not an artifact of the 7-market system's dimensionality: it emerges cleanly from the simplest possible specification.

5.5 Frequency Sensitivity of Information Shares

A natural concern is whether Polymarket's high-frequency dominance reflects informational leadership or merely its faster price-updating mechanism. Table 2 documents that update frequency varies by more than an order of magnitude across platforms, from 5.4% (Betfair) to 0.1% (Everygame). When a market carries forward a stale price, the VECM residual for that market is mechanically close to zero, which can inflate its estimated contribution to the permanent component—a well-known concern in the Hasbrouck IS literature (Baillie et al., 2002; De Jong, 2002).

If price staleness were the primary driver of information shares, then re-estimating at lower frequencies, where more markets have time to update, should compress the IS distribution and erode Polymarket's lead. Table 9 reports the results of this test, which is a direct evaluation of H4.

Table 9: Hasbrouck IS Midpoints by Sampling Frequency

Market	5-min	15-min	30-min	1-hour
Polymarket	0.429	0.278	0.192	0.119
Betfair	0.282	0.418	0.476	0.432
Bovada	0.028	0.054	0.046	0.061
BetOnline	0.046	0.041	0.043	0.041
William Hill	0.032	0.039	0.038	0.051
Unibet	0.101	0.117	0.108	0.086
Everygame	0.208	0.175	0.230	0.376

Notes: Each column re-estimates the 7-market VECM at the indicated sampling frequency. Lag orders adjusted to maintain similar time horizons. Hasbrouck IS midpoints are averages of upper and lower bounds across Cholesky orderings; for a given ordering, shares sum to one, but midpoints (which average across orderings) need not, and in these panels exceed one by 12–16% as a consequence. The ranking across venues, rather than the absolute magnitude, is the informative object.

Polymarket's IS declines monotonically from 43% at 5-minute frequency to 12% at 1-hour, while Betfair's *rises* from 28% to 43%—a clean crossover. We read this as the central evidence for H4, but we resist a structural interpretation. Because both venues are order books, the crossover *cannot* be attributed to an AMM-versus-limit-order-book distinction; it must instead reflect some difference in how the two books are used. Several non-exclusive mechanisms fit the pattern: Polymarket's crypto-native, 24/7 clientele may react first to news with aggressive marketable orders, so its innovations lead at the finest horizon, while Betfair's deeper, more experienced book may absorb the larger cumulative informed flow that comes to dominate the permanent

component once enough time has elapsed; differences in tick size, fees, and quote freshness could contribute as well. Our data cannot separate these channels. What we *can* say is that the crossover is not a pure speed-or-staleness artifact—a merely faster-updating venue would lead at all horizons, not surrender its lead as the horizon lengthens—and that the two exchanges appear to play complementary roles across horizons rather than one simply dominating.

Sportsbook IS remains low at all frequencies for the sharp books and William Hill (3–6%), confirming that their subordinate role is not merely a staleness artifact but reflects a lack of informational contribution. This is consistent with Proposition 3: bookmakers track the efficient price set elsewhere and contribute little innovation variance at any horizon. The exception is Everygame, whose IS *rises* from 21% at 5-minute to 38% at 1-hour frequency—the opposite of what the staleness explanation predicts. This anomaly likely reflects Everygame's large, discrete catch-up jumps, which at lower frequencies align with the estimation window and are mechanically amplified in the innovation variance. This pattern reinforces the decision to exclude Everygame from the clean panel.

5.6 Microstructure Mechanisms

Cross-market spread compression after Kalshi. At each 5-minute interval, we compute the spread between the highest and lowest implied probability across all active platforms. The mean cross-market spread fell from 7.2 pp pre-Kalshi to 4.2 pp post-Kalshi (October 4–November 5), a 42% reduction, and the 95th percentile compressed from 10.4 pp to 6.1 pp. To condition out confounds, we regress the daily max-min spread on a Kalshi entry dummy, realized volatility, the absolute distance from $p = 0.5$ (which mechanically compresses spreads), day-of-week and event fixed effects, using Newey–West HAC standard errors. The Kalshi dummy coefficient is -3.2 pp ($t = -7.78$, $p < 0.001$; Appendix Table A.13), indicating that cross-market spreads compressed significantly even after controlling for volatility and the approach of Election Day. Realized volatility enters positively ($t = 1.93$, $p = 0.054$), consistent with information-driven widening. Note, however, that the regression identifies conditional correlation rather than causation, as Kalshi's entry is confounded with election proximity.

Cross-correlogram evidence. Cross-correlograms between Polymarket and the six other common-panel markets at 5-minute frequency (Figure 4) are flat, with all pairwise correlations between 0.87 and 0.95 across lags. Contemporaneous correlations of 5-minute probability *changes* (Figure 5) are near zero (0.002–0.051), despite the very high correlation in levels. This contrast confirms that price discovery operates through the error correction channel rather than through simultaneous innovation, motivating the VECM approach.

5.7 Cross-Asset Evidence: Prediction Markets and Financial Markets

The within-prediction-market comparison cannot, by itself, establish that these venues perform a socially useful price-discovery function rather than merely racing one another to copy a common signal. We therefore step outside the prediction-market ecosystem and benchmark the Polymarket consensus against two liquid financial assets that price the same election outcome with exchange-grade timestamps: front-month E-mini S&P 500 futures and the equity of Trump Media & Technology Group (DJT). This exercise also speaks to the

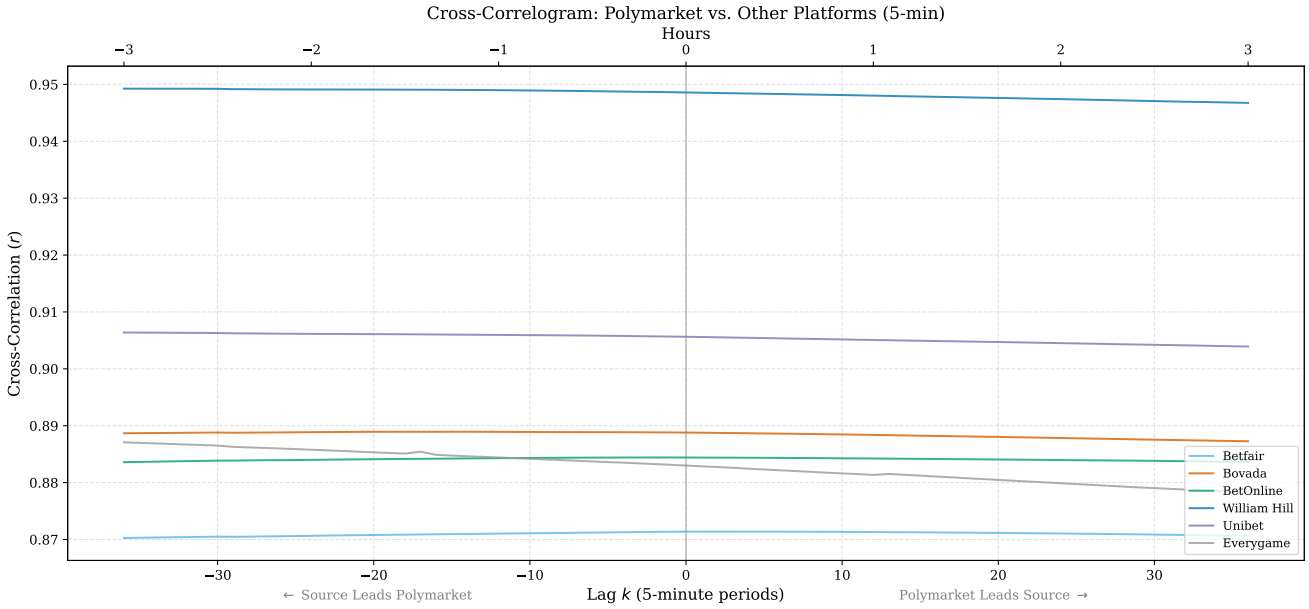


Figure 4: Cross-correlogram between Polymarket and six other platforms at 5-minute frequency. Positive lags indicate Polymarket leads.

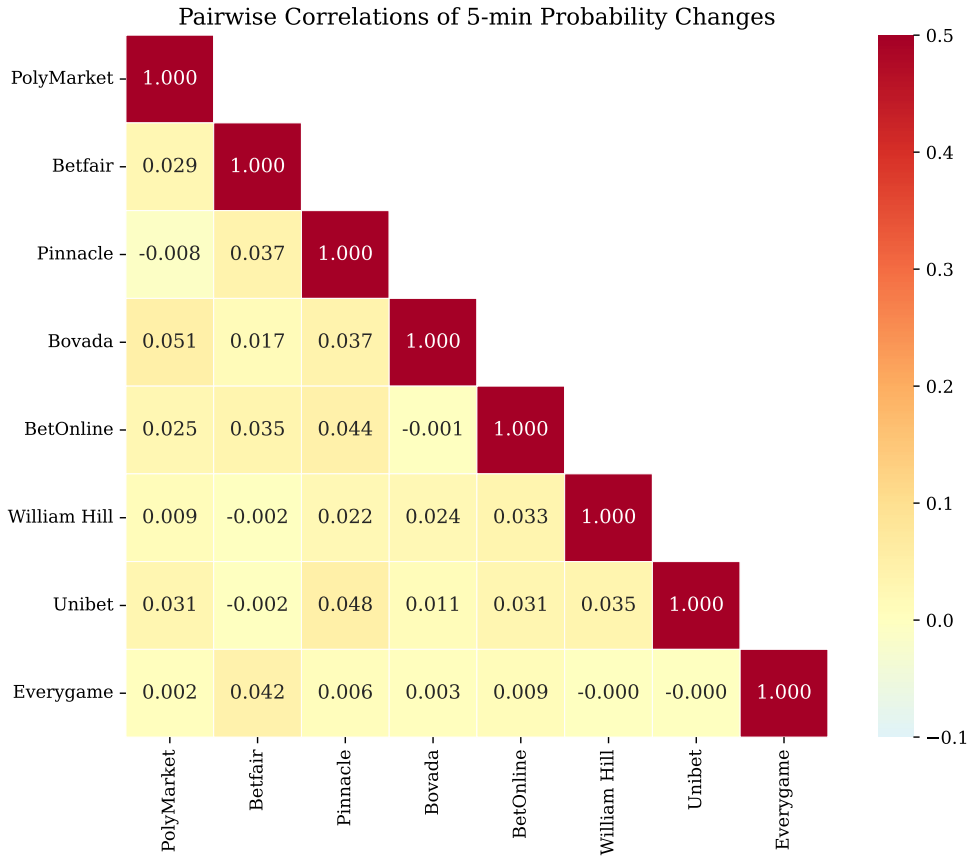


Figure 5: Pairwise correlations of 5-minute probability changes across all eight markets in the full panel (including Pinnacle from July onward; Kalshi excluded due to its shorter sample).

concern that our headline gap is mechanical. Financial markets are deep, attention-rich, and professionally arbitraged; a lead of the prediction market *over them* cannot be dismissed as one venue simply trying harder than another.

Election night. Figure 6 overlays the Polymarket probability and the S&P 500 future from the evening of November 5 into the following morning (UTC). The two series are strongly co-trending as the result resolves (level correlation 0.89), but the prediction market moves first at every stage. Because both repricings are gradual, the cleanest read is the timing of the cumulative move (Table 10, Panel A): Polymarket reaches 25% of its overnight move 15 minutes before the future, 50% twenty-five minutes before, and 75% and 90% roughly three to three-and-three-quarter *hours* before. Polymarket was 90% of the way to certainty by 04:50 UTC (11:50pm ET); the future did not complete the corresponding 90% of its rally until 08:35 UTC (3:35am ET). The 5-minute return cross-correlation is too noisy on so gradual and asynchronous a move to resolve the lead statistically, and the two series are not cointegrated within this twelve-hour window, so we report no within-window information share; the milestone timing is the robust statement, and it agrees with Flynn and Tarkom (2025): betting markets incorporate political information ahead of financial markets.

Weekend shocks. The advantage is starkest when financial markets cannot trade at all. Both of the campaign's exogenous weekend shocks—the July 13 assassination attempt (a Saturday) and the July 21 Biden withdrawal (a Sunday)—struck while the CME was closed for its 49-hour weekend break (Figure 7). Polymarket, trading 24/7, priced them continuously: the Trump probability moved about 11.5 percentage points over the assassination-attempt closure alone, before the S&P future could react. By the time futures reopened on Sunday evening, the assassination's residual impact on the index was negligible (a reopening gap of +0.07%); the political information had been impounded in the interim (Table 10, Panel B). For news that arrives on political rather than market time, 24/7 prediction markets fill a price-discovery role that financial markets structurally cannot.

DJT equity. Over the campaign's final months, the Polymarket probability and DJT—a near pure-play on a Trump victory—are cointegrated (Engle–Granger $p < 0.01$) and Granger-cause each other bidirectionally at 5-minute frequency during market hours. Within the trading session the two are tightly and roughly symmetrically linked: their Putniņš information leadership shares are both near one (1.02 for Polymarket, 0.99 for DJT), so neither is the clear intraday leader, and DJT's larger raw information share merely reflects its far greater idiosyncratic volatility. The direction of information flow is revealed instead by the *overnight* gap, when Polymarket trades but the stock does not. Regressing DJT's overnight return (previous close to next open) on Polymarket's contemporaneous overnight change in log-odds gives a strongly positive, significant loading ($\beta = 0.36$, $t = 3.5$, $R^2 = 0.22$, $N = 45$ sessions): information accumulating on Polymarket while the equity market is closed is impounded into DJT at the open. As with the weekend shocks, the prediction market leads the financial asset precisely when it alone can trade—an out-of-hours counterpart to Flynn and Tarkom (2025).

Taken together, the cross-asset results offer design-independent evidence for the paper's central claim: exchange-organized prediction markets perform genuine price discovery on political information, leading even deep, professionally traded financial markets around the events that defined the campaign.

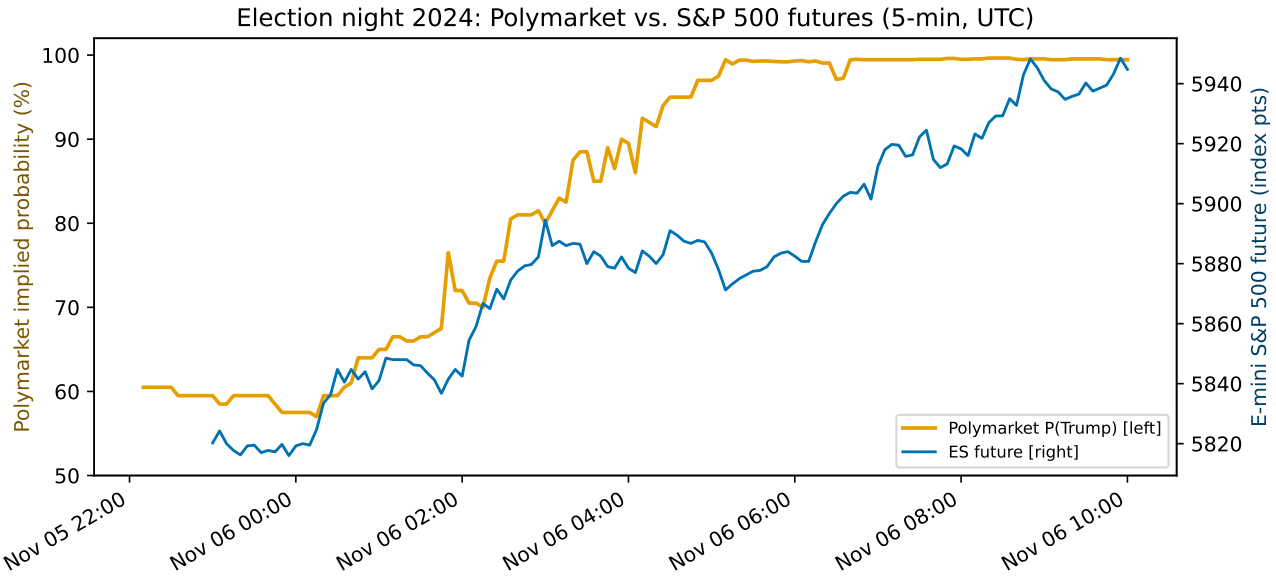


Figure 6: Election night 2024: the Polymarket Trump-victory probability (left axis) and the front-month E-mini S&P 500 future (right axis) at 5-minute frequency, UTC. The prediction market leads the future at every stage of the overnight repricing (Table 10, Panel A).

Table 10: Cross-Asset Price Discovery: Prediction Market vs. Financial Markets

Panel A: Election-night cumulative-move timing (UTC)

Milestone	Polymarket reaches	ES future reaches	ES lag (min)
25% of move	01:50	02:05	15
50% of move	02:35	03:00	25
75% of move	03:55	07:05	190
90% of move	04:50	08:35	225

Panel B: Weekend shocks (CME futures closed)

Event	Futures closed (h)	Polymarket move while closed (pp)	ES reopen gap (%)
Assassination attempt (Jul 13)	49	+11.5	+0.07
Biden withdrawal (Jul 21)	49	−3.5	+0.12

Notes: Panel A reports the time at which each series first reaches the indicated fraction of its total election-night move (Nov 5 22:00 to Nov 6 10:00 UTC), at 5-minute frequency; "ES lag" is the delay of the S&P 500 future relative to Polymarket. Panel B reports the duration of the CME weekend closure spanning each shock, the change in the Polymarket probability over that closure, and the S&P 500 future's gap at the Sunday-evening reopen. Financial series are from DataBento (E-mini S&P 500 futures, CME Globex MDP 3.0).

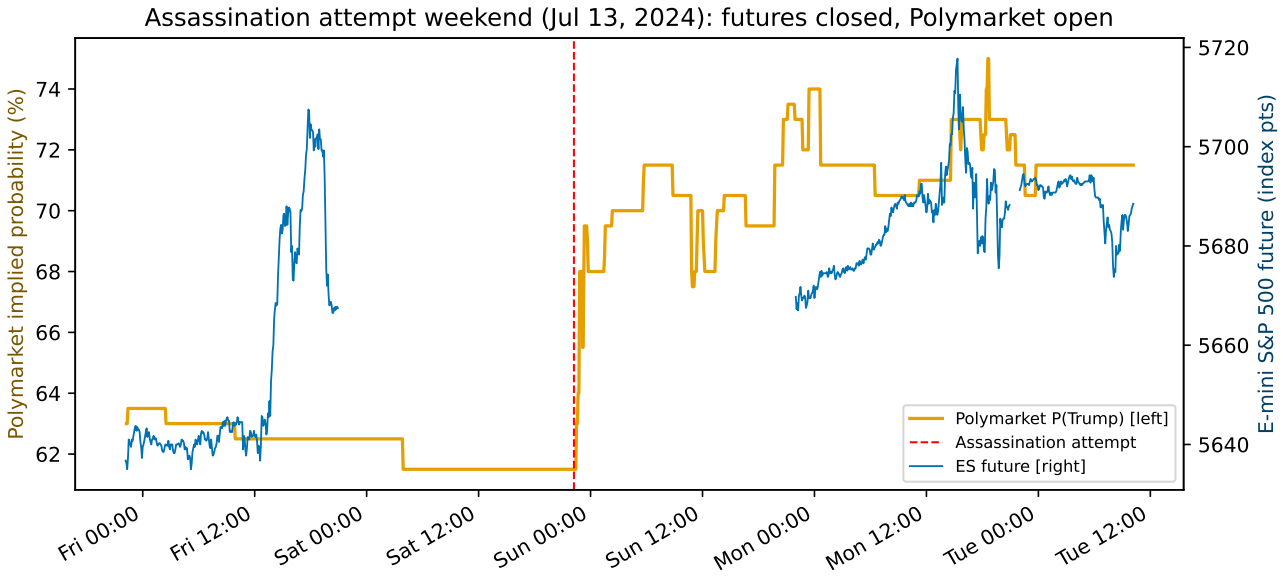


Figure 7: The July 13, 2024 assassination-attempt weekend. Polymarket (left axis) prices the shock continuously, while the E-mini S&P 500 future (right axis) is closed for the CME weekend break and cannot react until the Sunday-evening reopen. The dashed line marks the shock.

5.8 Robustness of the Ranking and Corroborating Evidence

Ordering invariance and weak exogeneity. The ordering-invariant Lien and Shrestha (2009) Modified Information Shares (MIS) closely track the Hasbrouck midpoints, confirming that the IS ranking is not driven by Cholesky ordering. In the Kalshi window, MIS similarly confirms Polymarket's leadership (47.4%) and Kalshi's nontrivial share (38.1%). Using MIS in the Putniņš (2013) ILS reproduces the split of Section 5.2: in the full sample every venue shows a speed advantage ($\text{ILS}_{\text{MIS}} < 1$; e.g. Polymarket 0.40, Betfair 0.11), while in the Kalshi window Polymarket (1.59) and Kalshi (1.57) flip to an informational advantage ($\text{ILS}_{\text{MIS}} > 1$) and the other venues remain below one. Weak exogeneity Wald tests provide a complementary perspective: all seven markets reject the null $H_0 : \alpha_{j,\cdot} = 0$ at the 1% level, meaning no venue is a pure price leader in the formal sense. However, the test statistics differ by an order of magnitude—Polymarket ($\chi^2 = 28.0$) and Betfair ($\chi^2 = 36.2$) versus BetOnline ($\chi^2 = 570$) and William Hill ($\chi^2 = 620$)—indicating far stronger statistical evidence of error correction for sportsbooks, consistent with their role as followers.

Connectedness, jumps, and breaks. Three further tools triangulate the same conclusions; we summarize them compactly rather than as separate findings. The Diebold and Yılmaz (2014) framework gives a total connectedness index of 55.9%. In the time domain Betfair is the strongest net transmitter (NET = +0.18) and Polymarket a net *receiver* (−0.21)—not a contradiction of Polymarket's high IS, but the same horizon-dependence as Table 9, since the index is computed at a 1-hour forecast horizon, exactly where Betfair leads. The Baruník and Křehlík (2018) frequency decomposition makes this explicit: the index splits into 2.5% short-run (<30 min), 8.2% medium-run, and 45.2% long-run (>2 hours), with Betfair's net transmission rising monotonically across bands, while in the Kalshi window both Polymarket and Kalshi transmit at every band—a within-model

version of the horizon crossover. The rolling index ranges from 4.1% (quiet May–June) to 75.2% (around the Biden withdrawal and the election), tracking information intensity. *Jumps.* The Lee and Mykland (2008) test finds 761 significant jumps; Polymarket and Betfair account for 91% of them, and of 45 multi-market jump clusters Polymarket leads 49% and Betfair 33%. No retail sportsbook ever leads a cluster, and the sharp book BetOnline (18%) is the only dealer venue that does. *Breaks.* Applying the Bai and Perron (1998) algorithm to the rolling IS series without imposing dates, 15 of 29 detected breaks (52%) fall within 14 days of a known political shock—the assassination attempt and Biden withdrawal sharpest among them—tying the regime shifts to identifiable news rather than to estimation noise.

The Model Confidence Set. Following Hansen et al. (2011), we ask whether a single venue can be statistically identified as the unique best price discoverer. We find that it cannot, and we report this as a cautionary result rather than a headline. Building the MCS loss panel from the full-sample bootstrap distribution of information shares (max-t statistic, 5,000 replications) yields a 90% MCS containing *only* Polymarket, with every other venue eliminated at $p < 0.02$. That procedure, however, treats each bootstrap draw as an independent observation and so understates the standard errors and overstates precision. When we instead build the loss panel from the rolling-window IS estimates—genuine, time-separated estimates whose dispersion reflects real estimation uncertainty—the 90% MCS contains *only Betfair*, and Polymarket is now excluded. The two exchanges trade places as the lone survivor depending on the specification, because Polymarket leads only late in the sample while Betfair leads on average across the campaign. This is exactly what the wide bootstrap confidence intervals of Section 5.2 imply: the data do not support crowning a single leader between Polymarket and Betfair. What *is* robust is the bottom of the ranking—Bovada, BetOnline, William Hill, and Unibet are excluded from the 90% MCS under *both* specifications. The MCS thus reinforces the exchange-over-dealer split while cautioning against any claim that one exchange is decisively the best.

6 Robustness

We examine the sensitivity of our findings across several dimensions. Detailed results are provided in the Appendix.

6.1 Specification Robustness

Including Everygame. Appendix Table A.1 reports the 7-market VECM including Everygame. Its artifactual IS of 21% is absorbed almost entirely by the two exchange venues: in the 7-market specification, Polymarket's IS falls from 47.5% to 42.9% and Betfair's from 37.1% to 28.2%. The exchange-over-sportsbook hierarchy is identical in both specifications; the clean panel simply provides a sharper estimate by removing the contaminated series.

Lags, rank, transformation, and sample. We verify that the IS ranking is robust to specification choices. Varying the lag order across $k \in \{13, 15, 17, 19, 21\}$ (BIC $\pm$ 4) yields Polymarket IS midpoints in the range 38.0–42.9%, with the venue ranking unchanged (Appendix Table A.3). Reducing the cointegrating rank from $r = 5$ to $r = 4$ preserves the exchange-over-sportsbook hierarchy (Appendix Table A.4); the BIC strongly favors $r = 1$ over $r = 5$, reflecting its heavy penalty for parameter-rich specifications, but since the economics dictate

a single common stochastic trend we maintain $r = n - 1$ and note that the IS ranking is insensitive to this choice. Replacing the logit transformation with raw probabilities produces near-identical results (Polymarket IS: 41.6% vs. 42.9% in the 7-market specification; Appendix Table A.5). Excluding overnight periods (8 PM to 8 AM ET) slightly increases Polymarket's IS, consistent with its 24/7 trading (Appendix Table A.6). Splitting the sample at July 22 (Appendix Table A.7), the two exchanges clearly lead in the high-information second half; in the quiet first half the ranking is more mixed—Betfair in particular is overtaken by several books—the same early-regime pattern visible in the rolling estimates (Figure 3). Pairwise Granger causality tests (Appendix Table A.8) confirm that Betfair Granger-causes all other markets ($p < 0.01$), with all 36 significant pairs surviving Holm–Bonferroni and Benjamini–Hochberg corrections at the 5% level.

6.2 Alternative Estimators and Inference

Bootstrap inference. The stationary bootstrap with BCa correction yields confidence intervals that are generally comparable to the fixed-block bootstrap, while providing two methodological improvements: (i) the optimal block length (22 periods $\approx$ 110 minutes) is data-driven rather than arbitrarily fixed at one day, and (ii) the BCa correction accounts for the skewness inherent in bounded $[0,1]$ IS estimates. The recursive-design wild bootstrap, which preserves the conditional heteroskedasticity structure of the original VECM residuals, produces similar confidence intervals, confirming that the GARCH effects present around political events do not materially distort the standard bootstrap inference.

DCC-GARCH time-varying shares. As an alternative to rolling-window VECMs with their arbitrary window-size choice, we fit a DCC-GARCH(1,1) model (Engle, 2002) to the VECM residuals and compute period-by-period IS using the time-varying covariance matrix $\mathbf{\Sigma}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$. The DCC parameters ($\hat{a} = 0.005$, $\hat{b} = 0.93$) indicate highly persistent but slowly evolving correlations. Polymarket's time-varying IS averages 31.6% with a 5th–95th percentile range of [20.4%, 71.2%], confirming the qualitative pattern from the rolling analysis while revealing finer-grained dynamics.

6.3 System and Staleness Robustness

Leave-one-out. To assess whether individual series drive the IS ranking, we re-estimate the VECM after dropping each of the seven markets in turn, computing 6-market IS for the remaining series (Appendix Table A.10). Polymarket's IS midpoint averages 41.2% across the six configurations in which it appears, ranging from 7.8% (when William Hill is dropped) to 59.2% (when Betfair is dropped). Betfair averages 47.5% across its six configurations. An exchange—Polymarket or Betfair—is the top venue in every one of the seven specifications, so the exchange lead does not hinge on any single series. The one disruption is the William Hill drop, where Polymarket's IS collapses to 7.8% and a sportsbook (Unibet, 9.9%) edges ahead of it for second place behind Betfair; this reflects the disturbance to the error-correction structure when a key slow-adjusting market is removed, not a decline in Polymarket's informational role.

Fresh-quote and Monte Carlo checks. The Everygame staleness artifact (Section 5.2) raises the question of whether forward-fill staleness inflates IS for other low-frequency markets. We address this with a fresh-quote filter: we identify 5-minute timestamps where the two most active markets (Polymarket and Betfair) have each

shown a non-zero return within the preceding 24 periods (2 hours). This retains 3,410 observations (20.7% of the sample)—a meaningful subset that excludes stale-price periods. On this fresh-quote subsample, Polymarket's IS is 40.1% [27.7%, 52.5%], consistent with its full-sample estimate (42.9%), confirming that exchange dominance is not an artifact of staleness.

As a complementary exercise, we simulate 1,000 synthetic cointegrated systems with equal true information shares and apply staleness at the empirically-calibrated update rates. Even with equal true shares, the *measured* shares diverge sharply once staleness is imposed: relative to the no-staleness benchmark, Betfair and the very stale Everygame gain measured IS (by +48.2 and +15.4 percentage points), while Polymarket loses a correspondingly large share. The lesson is not the specific reallocation but its magnitude—measured IS is highly sensitive to the staleness structure—which is why we rely on the clean-panel and fresh-quote estimates and read Everygame's inflated full-sample IS as an artifact.

6.4 Kalshi Window Sensitivity

The 32-day Kalshi window during peak election intensity warrants additional scrutiny. We draw 200 random 32-day windows from the pre-Kalshi period (January–September 2024) and estimate 7-market VECM(3, $r = 6$) for each, building a placebo distribution of IS midpoints (Appendix Table A.11). Polymarket's Kalshi-window IS in the 7-market specification (70.4%) falls at the 100th percentile of the placebo distribution (placebo mean: 27.7%, SD: 10.8%), far above the 95th percentile (44.3%). Sportsbooks' Kalshi-window IS values fall at or near the 0th percentile, suggesting the concentration of price discovery in exchanges is unusually high during this period even by the standards of similarly-sized windows. We also extend the frequency-sensitivity analysis to the 9-market Kalshi panel (Appendix Table A.12): Kalshi's IS persists across frequencies (40.9% at 5-minute, 31.2% at 15-minute, 33.9% at 30-minute, 25.8% at 1-hour), declining at lower frequencies but remaining among the top two venues at all horizons.

6.5 IS–Volatility Interaction

To test H3 formally—whether exchange IS increases disproportionately during high-information periods—we regress rolling IS on standardized realized volatility interacted with architecture dummies (DEX, Exchange, Sharp; Retail omitted), with market fixed effects and cluster-robust standard errors. None of the interaction coefficients are statistically significant (Vol×DEX: $t = -0.07$, $p = 0.95$; Vol×Exchange: $t = 0.30$, $p = 0.76$; Vol×Sharp: $t = 0.54$, $p = 0.59$). Market-by-market Newey–West regressions confirm the null: no venue shows a significant positive relationship between IS and volatility. This null result is informative: the exchanges' lead in price discovery is not concentrated in high-volatility episodes but is a persistent feature of the cross-section.

7 Conclusion

In a market with no clearing, no common regulator, and no arbitrageur to enforce the law of one price across venues, we asked where the efficient probability of a Trump victory is produced, and whether that production matters beyond the betting ecosystem. Applying the standard price-discovery toolkit to nine venues over the

2024 campaign—and benchmarking the resulting consensus against financial markets—we reach four conclusions, which we read descriptively.

First, the order-book exchanges lead price formation. Polymarket and Betfair jointly account for the vast majority of information shares, and every other measure—weak exogeneity, connectedness, jump timing, and the Hansen et al. (2011) Model Confidence Set, which excludes the dealer books under both constructions—places the sportsbooks in a following role. The split is robust to lag choice, cointegrating rank, transformation, and bootstrap method, and it characterizes the full sample and the high-salience final months, though the quiet early campaign is more mixed. We are careful about what this identifies: because election contracts are a sideline for the sportsbooks, the gap conflates venue organization with the attention a peripheral product attracts. One fact resists the pure-neglect reading—Pinnacle, the sharpest book in the market, takes political action seriously yet still follows—but we do not lean on a single venue. The honest statement is that, in this fragmented market, the exchanges discover prices and the dealer books follow, not that market design alone is responsible.

Second, leadership among the exchanges is horizon-dependent: Polymarket leads at the 5-minute frequency and Betfair at the 1-hour frequency, a crossover that appears both in frequency re-sampling and in the Baruník and Křehlík (2018) spectral decomposition. We document this pattern but interpret it cautiously. Both venues are order books, so the crossover is not an AMM-versus-limit-order-book phenomenon; it reflects some combination of clientele, depth, fees, and quote freshness that we cannot disentangle. What it does rule out is a pure speed-or-staleness story, under which a single venue would lead at all horizons.

Third, Kalshi's rapid attainment of nontrivial information shares after its October 2024 launch illustrates that a newly opened, regulated exchange can quickly gain informational presence. Wide confidence intervals preclude any claim of parity with Polymarket, and the 32-day window at peak salience precludes causal inference; we therefore read the Kalshi episode as a suggestive description rather than as evidence on the causal role of regulation.

Fourth, and most consequentially for the regulatory debate, the prediction-market consensus performs genuine price discovery relative to financial markets. It led S&P 500 futures throughout election night—by minutes early in the resolution and by hours late in it—priced the campaign's two weekend shocks while futures markets were closed for days, and led Trump Media's equity across the overnight gap when the prediction market trades but the stock does not. Because financial markets are deep and professionally arbitraged, this lead is the one piece of our evidence that cannot be explained by a venue "trying harder" than its prediction-market rivals. For regulators weighing whether event-contract exchanges serve a legitimate economic function, it is the most direct evidence we offer that they do.

Several limitations remain. We lack contract-level volume, so we cannot construct volume-weighted shares or test the liquidity channel directly; the nine series are heterogeneous price objects; and our cross-section is a single election. We have accordingly framed the within-ecosystem results as descriptive correlations rather than causal estimates of an architectural effect—though four patterns jointly narrow the alternatives: the horizon crossover, Kalshi's rapid entry on thin volume, the null IS–volatility interaction, and the architecture-grouped decomposition that reproduces the hierarchy in the most parsimonious specification. The natural next steps are transaction- and order-book-level data, as in Ng et al. (2025), and many-event settings such as

sports, where exchanges and sportsbooks compete in their core business and the comparison can be made on an equal footing. Still, the central fact needs no further data to be striking: a set of venues with no clearing, no shared regulator, and no arbitrageur between them nonetheless produced a probability that deep, professionally traded financial markets followed—the strongest case we can offer that exchange-organized event contracts perform a price-discovery function worth taking seriously.

References

Aktuğ, E. and Torul, O. (2026). Political capital: Real-time electoral probabilities and the value of corporate political alignment. SSRN Working Paper No: 5234903.

Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., Levmore, S., Litan, R., Milgrom, P., Nelson, F. D., Neumann, G. R., Ottaviani, M., Schelling, T. C., Shiller, R. J., Smith, V. L., Snowberg, E., Sunstein, C. R., Tetlock, P. C., Tetlock, P. E., Varian, H. R., Wolfers, J., and Zitzewitz, E. (2008). The promise of prediction markets. *Science*, 320(5878):877–878.

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):47–78.

Baillie, R. T., Booth, G. G., Tse, Y., and Zabotina, T. (2002). Price discovery and common factor models. *Journal of Financial Markets*, 5(3):309–321.

Baruník, J. and Křehlík, T. (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, 16(2):271–296.

Berg, J. E., Forsythe, R., Nelson, F. D., and Rietz, T. A. (2008). Results from a dozen years of election futures markets research. In *Handbook of Experimental Economics Results*, volume 1, pages 742–751. Elsevier.

Cavaliere, G. and Taylor, A. M. R. (2008). Bootstrap unit root tests for time series with nonstationary volatility. *Econometric Theory*, 24(1):43–71.

Chakravarty, S., Gulen, H., and Mayhew, S. (2004). Informed trading in stock and option markets. *The Journal of Finance*, 59(3):1235–1257.

Croxson, K. and Reade, J. J. (2014). Information and efficiency: Goal arrival in soccer betting. *The Economic Journal*, 124(575):62–91.

Cutting, L. E., Hughes-Berheim, S. S., Johnson, P. M., Baroud, H., and Goldstein, B. (2025). Are betting markets better than polling in predicting political elections? arXiv:2507.08921.

De Jong, F. (2002). Measures of contributions to price discovery: A comparison. *Journal of Financial Markets*, 5(3):323–327.

Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366):427–431.

Diebold, F. X. and Yılmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1):57–66.

Diebold, F. X. and Yılmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Business & Economic Statistics*, 32(1):128–146.

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman and Hall/CRC.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350.

Flynn, M. and Tarkom, A. (2025). How do financial markets price political uncertainty? evidence from the 2024 united states presidential election. *Finance Research Letters*, 75:106879.

Foucault, T. and Menkveld, A. J. (2008). Competition for order flow and smart order routing systems. *The Journal of Finance*, 63(1):119–158.

Gonçalves, S. and Kilian, L. (2004). Bootstrapping autoregressions with conditional heteroskedasticity of unknown form. *Journal of Business & Economic Statistics*, 22(1):89–101.

Gonzalo, J. and Granger, C. W. J. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, 13(1):27–35.

Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2):453–497.

Hanson, R., Oprea, R., and Porter, D. (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4):449–461.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *The journal of Finance*, 50(4):1175–1199.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica*, 59(6):1551–1580.

Lee, H. C. B., Stallaert, J., and Fan, M. (2020). Anomalies in probability estimates for event forecasting on prediction markets. *Production and Operations Management*, 29(2).

Lee, S. S. and Mykland, P. A. (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. *The Review of Financial Studies*, 21(6):2535–2563.

Lehar, A. and Parlour, C. A. (2021). Decentralized exchanges. *The Review of Financial Studies*, 34(11):5275–5313.

Lehmann, B. N. (2002). Some desiderata for the measurement of price discovery across markets. *Journal of Financial Markets*, 5(3):259–276.

Lien, D. and Shrestha, K. (2009). A new information share measure. *Journal of Futures Markets*, 29(4):377–395.

Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3(3):205–258.

Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *The Journal of Finance*, 75(2):585–617.

Manski, C. F. (2006). Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429.

Ng, H., Peng, L., Tao, Y., and Zhou, D. (2025). Price discovery and trading in modern prediction markets. SSRN Working Paper No. 5331995.

O'Hara, M. and Ye, M. (2011). Is market fragmentation harming market quality? *The Review of Financial Studies*, 24(11):3437–3483.

Page, L. and Clemen, R. T. (2013). Do prediction markets produce well-calibrated probability forecasts? *The Economic Journal*, 123(568):491–513.

Patton, A., Politis, D. N., and White, H. (2009). Correction to "Automatic block-length selection for the dependent bootstrap". *Econometric Reviews*, 28(4):372–375.

Pesaran, M. H. and Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1):17–29.

Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.

Politis, D. N. and White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23(1):53–70.

Putniņš, T. J. (2013). What do measures of price discovery tell us? an empirical investigation. *Journal of Financial Markets*, 16(4):619–648.

Rothschild, D. (2009). Forecasting elections: Comparing prediction markets, polls, and their biases. *Public Opinion Quarterly*, 73(5):895–916.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *The Journal of Finance*, 52(1):35–55.

Snowberg, E., Wolfers, J., and Zitzewitz, E. (2007). Partisan impacts on the economy: evidence from prediction markets and close elections. *The Quarterly Journal of Economics*, 122(2):807–829.

Snowberg, E., Wolfers, J., and Zitzewitz, E. (2013). Prediction markets for economic forecasting. In Elliott, G. and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 2. Elsevier.

Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126.

Yang, Z. and Tsang, K. P. (2026). The anatomy of a blockchain prediction market: Polymarket in the 2024 U.S. presidential election. SSRN Working Paper No. 6336679.

Appendices

A Robustness Tables

This appendix collects the supporting tables referenced in Sections 5 and 6.

Table A.1: 7-Market VECM Including Everygame (for Comparison)

	Hasbrouck IS			Bootstrap 90% CI			
Market	Lower	Mid	Upper	5^{th}	95^{th}	CS	Arch.
Polymarket	0.289	0.429	0.569	0.018	0.844	1.061	DEX
Betfair	0.125	0.282	0.439	0.000	0.303	1.171	Exchange
Bovada	0.000	0.028	0.056	0.000	0.295	0.006	Sharp
BetOnline	0.000	0.046	0.093	0.001	0.398	0.295	Sharp
William Hill	0.000	0.032	0.063	0.000	0.229	−0.007	Retail
Unibet	0.000	0.101	0.201	0.000	0.169	−0.841	Retail
Everygame	0.100	0.208	0.315	0.001	0.435	−0.686	Retail

Notes: 7-market VECM including Everygame. Everygame's 21% IS is an artifact of staleness (0.1% update rate); its residuals exhibit serial correlation (Ljung-Box $p < 0.01$). The clean 6-market panel (Table 4) is the primary specification. Bootstrap CIs from 500 block bootstrap replications. $k = 17$ lags, $r = 6$.

Table A.2: Block Bootstrap Confidence Intervals: Kalshi Window

	Hasbrouck IS (Bootstrap Distribution)		
Market	Mean	5% Lower	95% Upper
Polymarket	0.259	0.013	0.648
Kalshi	0.157	0.008	0.417
Betfair	0.112	0.001	0.411
Everygame	0.103	0.000	0.420
Unibet	0.085	0.000	0.422
Pinnacle	0.075	0.000	0.367
Bovada	0.074	0.000	0.361
William Hill	0.072	0.000	0.274
BetOnline	0.063	0.000	0.259

Notes: Results from 500 block bootstrap replications (1-day block size) on the 9-market panel during the 32-day Kalshi window. The wide and overlapping 90% confidence intervals for Polymarket and Kalshi reflect the short estimation window, confirming that the precise degree of informational parity is sensitive to sampling variation.

Table A.3: Lag Sensitivity: Hasbrouck IS Midpoints and GG Component Shares

	$k = 13$	$k = 15$	$k = 17$ (BIC)	$k = 19$	$k = 21$
Hasbrouck IS Midpoint					
Polymarket	0.401	0.380	0.429	0.427	0.428
Betfair	0.320	0.343	0.282	0.289	0.260
Bovada	0.026	0.027	0.028	0.026	0.033
BetOnline	0.064	0.052	0.046	0.064	0.081
William Hill	0.032	0.031	0.032	0.028	0.042
Unibet	0.102	0.100	0.101	0.090	0.094
Everygame	0.169	0.177	0.208	0.195	0.193
GG Component Share					
Polymarket	0.842	0.842	1.061	0.931	0.947
Betfair	1.022	1.095	1.171	1.032	0.989
Bovada	−0.057	−0.057	0.006	−0.030	−0.083
BetOnline	0.369	0.274	0.295	0.404	0.555
William Hill	−0.094	−0.014	−0.007	−0.054	−0.166
Unibet	−0.727	−0.767	−0.841	−0.683	−0.652
Everygame	−0.354	−0.373	−0.686	−0.600	−0.590

Notes: BIC selects $k = 17$ for the 7-market common panel. Cointegrating rank $r = 6$ in all specifications.

Table A.4: Cointegrating Rank Sensitivity: Hasbrouck IS Midpoints

	$r = 6$ $(n-1)$	$r = 5$ $(n-2)$	$r = 4$ $(n-3)$
Polymarket	0.429	0.382	0.611
Betfair	0.282	0.409	0.164
Bovada	0.028	0.050	0.038
BetOnline	0.046	0.037	0.030
William Hill	0.032	0.039	0.126
Unibet	0.101	0.077	0.086
Everygame	0.208	0.104	0.109

Notes: 7-market common panel with $k = 17$ lags. The theoretical rank is $n - 1 = 6$ (one common stochastic trend).

Table A.5: Logit vs. Raw Probability: Hasbrouck IS Midpoints

	Common Panel		Kalshi Window	
Market	Logit	Raw	Logit	Raw
Polymarket	0.429	0.416	0.438	0.424
Betfair	0.282	0.286	0.057	0.062
Bovada	0.028	0.030	0.018	0.019
BetOnline	0.046	0.056	0.002	0.001
William Hill	0.032	0.035	0.002	0.002
Unibet	0.101	0.111	0.000	0.000
Everygame	0.208	0.202	0.048	0.049
Pinnacle	—	—	0.031	0.032
Kalshi	—	—	0.409	0.417

Notes: "Logit" columns use the baseline logit transformation $\log(p/(1-p))$. "Raw" columns use untransformed implied probabilities. VECM specification otherwise identical.

Table A.6: Overnight Exclusion: Hasbrouck IS and GG Component Shares

	Full 24-hour		Excl. 8PM–8AM ET	
Market	IS Mid	CS	IS Mid	CS
Polymarket	0.429	1.061	0.439	1.504
Betfair	0.282	1.171	0.173	1.280
Bovada	0.028	0.006	0.051	0.480
BetOnline	0.046	0.295	0.055	0.482
William Hill	0.032	−0.007	0.062	−0.435
Unibet	0.101	−0.841	0.109	−1.134
Everygame	0.208	−0.686	0.297	−1.178

Notes: Excluding overnight hours (8PM–8AM Eastern) removes periods when U.S. sportsbooks are least active.

Table A.7: Sub-Sample Stability: Pre-July vs. Post-July

	Jan–Jul 22		Jul 22–Nov	
Market	IS Mid	CS	IS Mid	CS
Polymarket	0.267	1.540	0.362	1.177
Betfair	0.116	0.275	0.268	1.083
Bovada	0.106	−0.429	0.047	0.116
BetOnline	0.254	−0.777	0.127	0.702
William Hill	0.266	1.106	0.069	−0.264
Unibet	0.068	−0.024	0.093	−0.514
Everygame	0.157	−0.691	0.265	−1.298

Notes: Sample split at July 22, 2024 (after both the assassination attempt and Biden's withdrawal).

Table A.8: Pairwise Granger Causality p-Values (5-Minute First Differences)

From ↓ / To →	PM	BF	BOV	BO	WH	UNI	EG
PM	—	0.000	0.000	0.455	0.510	0.600	0.000
BF	0.000	—	0.000	0.000	0.000	0.000	0.000
BOV	0.000	0.000	—	0.000	0.000	0.000	0.000
BO	0.000	0.000	0.000	—	0.000	0.000	0.000
WH	0.000	0.000	0.000	0.720	—	0.000	0.000
UNI	0.000	0.000	0.000	0.005	0.354	—	0.000
EG	0.000	0.000	0.103	0.000	0.000	0.000	—

Notes: Entry (i, j) is the p-value for the null hypothesis that market i does not Granger-cause market j. Tests conducted on 5-minute first differences of logit-transformed probabilities.

Table A.9: Architecture-Grouped IS Decomposition

Specification	Group	IS Lower	IS Mid	IS Upper	GG Share
3-Group	DEX	0.522	**0.642**	0.762	0.662
	Exchange	0.210	0.330	0.450	0.581
	Sportsbook	0.001	0.051	0.102	−0.243
3-Group Clean	DEX	0.479	**0.626**	0.773	0.541
	Exchange	0.185	0.330	0.475	0.451
	Sportsbook	0.000	0.133	0.266	0.009
4-Group	DEX	0.369	**0.525**	0.681	0.677
	Exchange	0.081	0.224	0.368	0.530
	Sharp	0.003	0.173	0.343	0.574
	Retail	0.013	0.162	0.310	−0.781

Notes: Markets averaged within architecture groups. 3-Group: DEX = Polymarket, Exchange = Betfair, Sportsbook = 5 sportsbooks. 3-Group Clean excludes Everygame. 4-Group: Sharp = Bovada + BetOnline, Retail = William Hill + Unibet + Everygame. VECM(5) with $r = n - 1$.

Table A.10: Leave-One-Out IS Sensitivity: Polymarket IS Midpoint When Each Market Is Dropped

Dropped Market	PM IS Mid	PM IS Lower	PM IS Upper	BF IS Mid	n
Betfair	0.592	0.470	0.714	—	6
Bovada	0.423	0.282	0.565	0.302	6
BetOnline	0.423	0.310	0.536	0.488	6
William Hill	0.078	0.025	0.132	0.784	6
Unibet	0.482	0.345	0.618	0.225	6
Everygame	0.475	0.375	0.576	0.371	6

Notes: Each row drops the indicated market and re-estimates a 6-market VECM(17, $r = 5$). Polymarket's IS averages 41.2% across configurations. The outlier when William Hill is dropped (7.8%) reflects disruption to the error correction structure.

Table A.11: Kalshi Placebo Windows: Actual vs. Pre-Kalshi Null Distribution

	Placebo Distribution ($N = 200$)					
Market	Mean	SD	5th pct	95th pct	Actual	Pct. rank
Polymarket	0.277	0.108	0.098	0.443	**0.704**	100.0%
Betfair	0.153	0.155	0.024	0.539	0.155	59.5%
Bovada	0.108	0.045	0.054	0.192	0.017	0.0%
BetOnline	0.132	0.068	0.024	0.203	0.061	24.5%
William Hill	0.204	0.152	0.048	0.402	0.001	0.0%
Unibet	0.106	0.043	0.049	0.165	0.006	0.0%
Everygame	0.218	0.105	0.092	0.405	0.062	0.5%

Notes: 200 random 32-day windows drawn from the pre-Kalshi period (Jan–Sep 2024). Each window estimated with 7-market VECM(3, $r = 6$). Percentile rank indicates the fraction of placebo IS values below the actual Kalshi-window IS.

Table A.12: Kalshi-Window Frequency Sensitivity

Market	5min	15min	30min	1h
Polymarket	0.438	0.538	0.455	0.018
Kalshi	0.409	0.312	0.339	0.258
Betfair	0.057	0.052	0.119	0.216
Pinnacle	0.031	0.061	0.015	0.054
Bovada	0.018	0.001	0.003	0.008
BetOnline	0.002	0.015	0.040	0.052
William Hill	0.002	0.010	0.074	0.003
Unibet	0.000	0.003	0.010	0.030
Everygame	0.048	0.022	0.004	0.414

Notes: IS midpoints for the 9-market Kalshi panel at different sampling frequencies. Kalshi's IS persists across frequencies, declining from 40.9% (5-minute) to 25.8% (1-hour) but remaining among the top two venues at all horizons. At the 1-hour frequency, Polymarket's IS drops to 1.8% while Everygame's rises to 41.4%, likely reflecting staleness-driven redistribution in this short, 32-day sample.

Table A.13: Cross-Market Spread Regression

Variable	Coefficient	NW SE	t-stat	p-value
Constant	0.066	0.006	10.40	< 0.001
Kalshi Dummy	-0.032	0.004	-7.78	< 0.001
Realized Vol	5.778	2.996	1.93	0.054
$\lvert p - 0.5 \rvert$	0.060	0.063	0.95	0.342
Monday	-0.011	0.005	-2.25	0.025
Election Day	-0.018	0.007	-2.43	0.015

Notes: OLS regression of daily max-min cross-market probability spread on covariates. Newey–West HAC standard errors. $N = 310$ trading days. Kalshi dummy = 1 after October 4, 2024.

B Conceptual Framework: Venue Choice and Price Discovery

This appendix presents a partial-equilibrium framework that rationalizes our cross-sectional findings and maps them to hypotheses H1 and H2. The framework is deliberately simple—it features no market clearing, no strategic interaction, and no welfare analysis—and the propositions below are transparent intuition-building sketches under strong simplifying assumptions, not general theorems. They are meant to organize the empirical patterns, not to deliver structural restrictions beyond those captured by our reduced-form hypotheses. We do not attempt to formalize H4 (horizon-dependent leadership): because the two leading venues are both order books, the candidate channels (clientele, depth, fees, quote freshness) cannot be separated in a tractable way, and we treat that pattern as an empirical observation in the main text.

B.1 Environment

Time is continuous on $[0, T]$. A single binary event $Y \in \{0,1\}$ is realized at time T with prior probability $p_0 = \mathbb{P}(Y = 1)$.

There is a continuum of risk-neutral traders. A measure μ_I are *informed* and receive private information about Y; the remaining traders generate uninformed (noise) order flow. Conditional on Y, each informed trader i receives a private signal s_i at random times according to a Poisson process with intensity $\lambda > 0$. Signals are independent across traders and satisfy the Monotone Likelihood Ratio Property so that each signal induces a posterior belief $p_s \in (0,1)$ with $p_s \neq p_0$ almost surely.

There are three trading venues indexed by $v \in \{A, E, B\}$:

1. An on-chain, order-driven exchange A (Polymarket-type): a thin but fast-reacting central limit order book.
2. A deeper order-driven exchange E (Betfair/Kalshi-type).
3. A bookmaker B (dealer sportsbook).

Each venue quotes a price $P_v(t) \in (0,1)$ for an Arrow–Debreu-style contract paying 1 if $Y = 1$ and 0 otherwise. We interpret $P_v(t)$ as the venue's implied probability of $Y = 1$.

B.2 Venue Technologies

Venues differ along two dimensions: depth (price impact) and participation cost. (The bookmaker additionally revises its price only at discrete review times, modeled in Section B.5.)

Depth and price impact. Trading q_v contracts at venue v when the current price is P_v moves the price by

$$\Delta P_v = \kappa_v q_v, \qquad \kappa_A > \kappa_E > \kappa_B \approx 0. \tag{6}$$

The on-chain exchange A has the thinnest book (highest price impact), the exchange E has a deeper book (lower impact), and the bookmaker absorbs order flow at (essentially) fixed odds except at review times.

Participation costs. A trader who trades q_v contracts at venue v incurs a per-unit cost $c_v \geq 0$, capturing access frictions, regulatory constraints, and perceived counterparty risk. We allow c_v to differ across trader

types. In particular, a mass $\mu_R \le \mu_I$ of "regulated" informed traders face $c_A = c_E^{\text{off}} = +\infty$ (they cannot use offshore or on-chain venues) but finite c_E^{reg} at a regulated exchange E; the remaining informed traders can use all venues.

B.3 Informed Trading at a Given Venue

Consider a trader with posterior belief p_s who arrives at time t when venue v quotes price $P_v(t)$. She chooses q_v to maximize expected profit net of price impact and participation cost:

$$\Pi_v(q_v) = q_v(p_s - P_v(t)) - \frac{1}{2}\kappa_v q_v^2 - c_v q_v. \tag{7}$$

The first term is expected payoff from the contract, the second term is the cost induced by price impact (6), and the third term is the participation cost.

Proposition 1 (Optimal order size and venue comparison)**.** *For given $(p_s, P_v, c_v, \kappa_v)$, the unique optimal order at venue v is*

$$q_v^* = \frac{p_s - P_v - c_v}{\kappa_v} \quad \text{if } p_s - P_v > c_v, \qquad q_v^* = 0 \text{ otherwise.}$$

The associated maximal profit is

$$\Pi_v^* = \frac{(p_s - P_v - c_v)^2}{2\kappa_v} \cdot \mathbb{1}\{p_s - P_v > c_v\}.$$

For a given signal realization p_s, an informed trader will submit her marginal order to the venue v that maximizes Π_v^* among those with finite c_v.

Proof. The objective $\Pi_v(q_v)$ is strictly concave in q_v for any $\kappa_v > 0$, with derivative

$$\frac{\partial \Pi_v}{\partial q_v} = p_s - P_v - \kappa_v q_v - c_v.$$

The first-order condition $\frac{\partial \Pi_v}{\partial q_v} = 0$ yields $q_v^* = (p_s - P_v - c_v)/\kappa_v$. If $p_s - P_v \le c_v$, the derivative at $q_v = 0$ is weakly negative and the optimal order is $q_v^* = 0$. Substituting q_v^* back into $\Pi_v(q_v)$ gives the expression for Π_v^*. Since informed traders are risk-neutral and trading opportunities across venues are independent, a trader allocates her marginal order to the venue with highest Π_v^* subject to feasibility (finite c_v). □

B.4 Regulated Entry and Information Leadership

We now extend the model to two exchange-type venues: an incumbent offshore exchange E_o and a newly opened regulated exchange E_r (Kalshi-type). Both share similar depth but serve different clienteles.

Before E_r opens, a mass μ_R of "regulated" informed traders face $c_A = c_{E_o} = +\infty$ and cannot trade; only the mass μ_U of unconstrained informed traders participate on A and E_o. At time $t = t^*$, E_r opens with participation cost $c_{E_r} < +\infty$ for the regulated traders and depth parameter $\kappa_{E_r} \approx \kappa_{E_o}$.

Proposition 2 (Rapid catch-up of a regulated exchange)**.** *Suppose that (i) $\mu_R > 0$ informed traders can only trade at the regulated exchange E_r after entry, (ii) κ_{E_r} is comparable to κ_{E_o}, and (iii) the arrival rate and informativeness of signals for regulated traders are similar to those of unconstrained traders. Then, over a sufficiently*

short post-entry window $[t^*, t^* + \Delta]$*, the regulated exchange* E_r *can contribute a share of efficient-price innovations comparable to that of the incumbent exchange* A*, even if its total trading volume is smaller.*

Sketch of argument. After t^*, each regulated trader with signal p_s solves the same problem as in Proposition 1 but can only choose E_r. Given κ_{E_r} and c_{E_r}, each signal generates a strictly positive optimal order $q^*_{E_r}$ whenever $|p_s - P_{E_r}| > c_{E_r}$. Since $\mu_R > 0$ and signals arrive at rate λ, the aggregate informed order flow to E_r over $[t^*, t^* + \Delta]$ is of order $\mu_R \lambda \Delta$, which can be large even if Δ is modest. Under the maintained assumption that the distribution of signals for regulated traders is as informative as for unconstrained traders, these orders are *highly selected* on information: they come from traders who were previously unable to trade anywhere.

By contrast, the incumbent exchange A draws informed flow only from the unconstrained traders, of measure μ_U. A convenient feature of the linear-impact model is that the price move induced by one informed order, $\kappa_v q^*_v = p_s - P_v - c_v$, is independent of the venue's depth κ_v: each signal moves the price by its own "surprise," regardless of where it is traded. The variance a venue contributes to the efficient-price innovation is therefore governed by the *number* of informed orders it attracts—of order $\mu\lambda\Delta$—times the signal-surprise variance. Since E_r draws the previously-excluded mass μ_R and A the mass μ_U, their contributions stand in rough proportion $\mu_R : \mu_U$. For μ_R not too small relative to μ_U, E_r's information share can quickly approach A's even though its *volume* remains lower. □

Proposition 2 provides a simple mechanism through which a newly launched regulated exchange can attain an information share comparable to an established prediction market's within a few weeks, as the point estimates for Kalshi in Table 5 suggest—though, as the main text emphasizes, the wide confidence intervals there caution against reading this as exact parity. This is the theoretical foundation for H2.

B.5 Bookmakers as Informational Followers

We finally sketch a simple model in which the bookmaker acts as an informational follower. Unlike A and E, the bookmaker does not adjust prices mechanically with order flow. Instead, it periodically sets odds to balance expected profit against the reputational or regulatory cost of visibly diverging from the "efficient" probability.

Let $\tilde{P}(t)$ denote the efficient price implied by the order-book exchanges A and E. At discrete review times $\{t_k\}_{k\geq 1}$, the bookmaker chooses a price $P_B(t_k)$ to solve

$$\max_{P_B} \mathbb{E}\big[\text{betting profit}(P_B) \,|\, \tilde{P}\big] \; - \; \gamma\big(P_B - \tilde{P}\big)^2, \qquad \gamma > 0. \tag{8}$$

The second term captures the cost of quoting odds that deviate from the consensus efficient price $\tilde{P}$.

Proposition 3 (Bookmakers track the efficient price)**.** *Suppose that (i) bettor demand is symmetric around the perceived efficient price* $\tilde{P}$ *so that expected betting profit is locally maximized at* $P_B = \tilde{P}$ *in the absence of the penalty term, and (ii) the penalty parameter* $\gamma > 0$*. Then, at each review time* t_k*, the bookmaker's optimal odds satisfy*

$$P^*_B(t_k) = \tilde{P}(t_k).$$

If reviews occur only when $|P_B - \tilde{P}|$ exceeds a threshold $\theta > 0$, then between reviews P_B is a piecewise-constant, lagged, and thresholded version of $\tilde{P}$.

Proof. By assumption (i), without the penalty term, the bookmaker's expected profit is locally maximized at $P_B = \tilde{P}$. A second-order Taylor expansion around $\tilde{P}$ implies that, to a first approximation, the expected profit term is quadratic and symmetric around $\tilde{P}$. Adding the quadratic penalty in (8), the overall objective is strictly concave and uniquely maximized at $P_B^* = \tilde{P}$. If the bookmaker only reviews odds when $|P_B - \tilde{P}| > \theta$, then P_B remains fixed between review times and jumps toward $\tilde{P}$ when the deviation hits the threshold. Thus, the bookmaker's odds process is a lagged and thresholded transformation of the efficient price. □

Proposition 3 implies that the bookmaker rarely initiates price changes and mostly adjusts odds in response to movements in the more informative venues. In a cointegrated system, this generates:

- small or zero information shares (the bookmaker contributes little to the variance of the efficient price innovation), and
- negative or near-zero component shares (bookmaker odds overshoot and then mean-revert toward the efficient price).

These predictions match our empirical findings (Table 4) and provide the theoretical foundation for H1.