Adverse Selection in Prediction Markets: Evidence from Kalshi

Robert Bartlett* Maureen O'Hara†

March 26, 2026

Abstract

Using 41.6 million trades on Kalshi, where the taker side is disclosed and every contract settles, we measure adverse selection in prediction markets without trade-classification error. Adapting Kyle's λ and the Glosten-Harris decomposition via a log-odds transformation, we find that single-name contracts (referencing specific companies or individuals) exhibit significantly higher informed price impact than broad-based contracts ($t = 10.88$). Effective spreads are also wider, yet makers earn more on single-name con tracts than broad-based contracts where adverse selection is more modest (3.94 vs. 0.74 cents per trade). The resolution is a frequency-magnitude decomposition of maker profits: single-name makers exploit a frequency edge (takers overbet YES on contracts that settle NO) that generates an optimism tax exceeding the cost of informed trading by a factor of 1.52×. Adapting VPIN to prediction markets as a real-time toxicity metric, we find that one-sided flow predicts maker losses on single-name contracts because its direction forecasts the settlement outcome. On broad-based contracts, where flow reflects the trajectory of a continuously moving underlying rather than private information about the terminal value, it does not. The favorite-longshot bias generates a behavioral surplus that cross-subsidizes adverse selection costs, sustaining liquidity even where informed trading is most severe.

JEL Codes: G10, G14, G23 *Keywords:* prediction markets, adverse selection, informed trading, market microstructure, event contracts

*Stanford Law School; Stanford Graduate School of Business. Email: rbartlett@law.stanford.edu.
†Cornell University, Johnson College of Business. Email: mo19@cornell.edu.

1 Introduction

On February 9, 2026, Kalshi (the first CFTC-regulated prediction market exchange) settled the contract "Will Jeff Bezos attend the Super Bowl?" at NO. Over the preceding week, 888,006 contracts had changed hands on this single question. The *Wall Street Journal* later identified the source of informed trading: Evan Whitesell, Bezos's stepson and a member of Sigma Alpha Epsilon at the University of Miami mentioned his family's Super Bowl plans at a fraternity event, and fraternity brothers traded on the information.[1] The YES price fell from roughly 70 cents to 30 cents as fraternity networks bought NO. This is, by any definition, informed trading: someone with private knowledge of the event outcome was on the other side of the book from the liquidity providers. Yet the makers on this contract still earned $29,274 in aggregate.

How? Of the 888,006 contracts traded, 518,000 were YES purchases by uninformed takers who evidently believed Bezos would attend. These trades were profitable for makers when Bezos did not attend, and the profits more than offset the losses that makers took against the informed fraternity traders. A companion contract on Mark Wahlberg's attendance makes the point even more starkly: a false rumor originating from a Clemson University fraternity generated 24 million contracts, takers were 68.6% YES-biased, and when Wahlberg did not attend, makers earned $1.18 million. In this case there were no informed traders at all; the makers' entire profit came from takers who mistook a rumor for a signal.

These episodes raise a question that is central to the economics of market making: can the adverse selection framework developed for equity markets (Glosten and Milgrom [1985], Kyle [1985], and the extensive empirical literature that followed [Glosten and Harris, 1988, Hasbrouck, 1991, Huang and Stoll, 1996]) be applied to prediction markets? And if informed trading is present, as the Bezos contract suggests, why do prediction market makers appear to do so well? The astonishing volume growth of Kalshi, particularly over the last year, underscores the importance of understanding these vibrant new markets (see Figure 1).

[1] "The Wildest Frat Party on Campus: Prediction Markets," *Wall Street Journal*, March 5, 2026.

The Bezos incident is not an isolated curiosity. Kalshi has sanctioned traders for exploiting nonpublic information on at least two occasions (a political candidate who bet on his own race and a media employee who traded on embargoed content), and the exchange's Super Bowl markets alone generated over $1 billion in volume amid widespread discussion of insider trading risk.[2] But Kalshi also lists hundreds of thousands of contracts on macroeconomic indicators, market indices, and asset prices where the risk of informed trading is far less acute. No individual possesses a decisive private signal about where the S&P 500 will close tomorrow or what the next CPI print will be. The juxtaposition of contracts with concentrated private information and contracts with dispersed public information, all trading on the same exchange through the same order book, creates a natural setting for studying how varying levels of adverse selection affect market making in prediction markets.

We exploit this variation using the trading record from Kalshi, drawing on 41.6 million trades across 478,167 contracts that we classify as either *single-name* (referencing a specific company or individual) or *broad-based* (referencing macroeconomic aggregates, market indices, or asset prices).[3] Kalshi's data offer two features that are, to our knowledge, unavailable in any equity market dataset. First, every trade record discloses the taker side (whether the party crossing the spread bought YES or NO), eliminating the need for trade-classification algorithms. Second, every contract settles at $1 (YES) or $0 (NO), so the realized profit or loss for the maker on every trade can be computed exactly.

The parallel to single-stock versus index-ETF trading in equity markets [Madhavan, 2002] is direct: concentrated information environments should produce more adverse selection.

[2] "Kalshi Found Some Insider Traders," Matt Levine, *Bloomberg Opinion*, February 25, 2026. See also "Predicting the Big Game," Matt Levine, *Bloomberg Opinion*, February 9, 2026; "Prediction markets allow trading on Super Bowl commercials, prompting insider trading questions," *CNBC*, January 30, 2026.

[3] Kalshi lists over 7 million markets spanning politics, sports, weather, entertainment, and other categories, though 87% are programmatically generated parlay contracts that bundle combinations of the roughly 1 million standalone markets into multi-leg bets. We restrict the main analysis to single-name and broad-based standalone contracts where the cross-sectional variation in information concentration should be especially acute (see Appendix A for details on how we classify Kalshi contracts). Appendix B tests this conjecture by estimating our adverse selection metrics across all non-parlay Kalshi contract categories. Single-name contracts exhibit the highest informed price impact of any category and broad-based among the lowest, with politics, sports, entertainment, weather, and other contracts falling in between.

Single-name and broad-based contracts trade on the same exchange, through the same order book, with the same population of takers and makers; only the information structure differs.

Adapting Kyle's λ to the prediction market setting using a log-odds transformation [Tsang and Yang, 2026], we find that they do. Single-name contracts exhibit significantly higher informed price impact ($t = 10.88$, $p < 0.001$). A Glosten-Harris decomposition [Glosten and Harris, 1988] confirms that this impact is almost entirely permanent (information, not inventory), with the transitory component economically negligible. These results hold with and without controls, when restricting to financial contracts only, and when using a continuous measure of information concentration rather than the binary single-name/broad-based classification.

Effective spreads are also wider on single-name contracts ($t = 3.17$), the conventional prediction. But in prediction markets, wider spreads do not protect makers the way they do in equities. The equity market maker earns the spread on the round trip; the prediction market maker holds a directional position to settlement. Her profit depends on the gap between the trade price and the settlement value, not on the bid-ask spread. Full collateralization makes offsetting trades capital-intensive rather than capital-releasing, so the spread is not a compensation mechanism in the traditional sense. The question is not whether spreads widen—they do—but what actually sustains maker profitability in the face of adverse selection.

The answer lies in how makers actually make money, and this is where the prediction market setting is genuinely illuminating.[4] Because we observe both the taker side and the settlement outcome for every trade, we can decompose maker profits into two components: a *frequency edge* (winning more often than losing) and a *magnitude edge* (winning bigger amounts). On single-name contracts, makers profit almost entirely through frequency: tak-

[4]Unlike equity markets, where market making is dominated by large institutional firms such as Citadel Securities or Susquehanna, Kalshi's maker population is fragmented. The exchange has noted that "only 5% of bid matches on the platform are from the big institutional market makers you think of," with the rest coming from "over 2,000 smaller market makers, some of which are just individuals" [eFinancialCareers, 2025]. The maker profitability results that follow should be understood in this context: the makers absorbing adverse selection on Kalshi are predominantly small, non-institutional participants.

ers are disproportionately YES-biased (60.2% buy YES), but only 30.9% of single-name contracts settle YES, so makers win on 56.6% of trades at roughly symmetric magnitudes (+29.24 cents per win, −29.06 cents per loss). The frequency edge accounts for 98% of maker profits on single-name contracts: +3.85 cents of a total +3.94 cents per trade.

Broad-based contracts, where takers are better calibrated, tell a starkly different story. Makers barely break even, earning just 0.74 cents per trade—over five times less than on single-name contracts. The frequency and magnitude edges are both small (+0.41 and +0.33 cents, respectively). Makers earn *far more* on single-name contracts, where adverse selection is most severe.

What makes this equilibrium work is the well-documented favorite-longshot bias—the tendency for bettors to overpay for low-probability outcomes [Wolfers and Zitzewitz, 2004]. On Kalshi, takers systematically overbet YES on contracts that settle NO, transferring wealth to makers. Becker [2025] and Bürgi, Deng, and Whelan [2026] have documented this "Optimism Tax" from the taker's perspective; our contribution is to show that the behavioral surplus is concentrated in single-name contracts, precisely where adverse selection is highest, and that it is large enough to cover the losses that informed traders impose on makers. The Bezos contract illustrates the mechanism: fraternity insiders imposed losses by buying NO, but the flood of uninformed YES takers more than compensated.

To investigate whether informed trading can be detected in real time, we adapt the VPIN flow toxicity metric [Easley, López de Prado, and O'Hara, 2012] to the prediction market setting. Standard equity VPIN requires several modifications for binary contract markets, including a restriction to interior prices (30–70 cents) to address a settlement-outcome endogeneity that has no equity-market analog: as a contract's outcome becomes clear and its price converges to 0 or 100, flow becomes mechanically one-sided, inflating VPIN without reflecting informed trading. At interior prices, where this confound is absent, trailing VPIN predicts maker losses on the next volume bucket for single-name contracts (OLS $t = -2.42$; logit $z = 10.56$) but not for broad-based (OLS $t = 0.56$). The divergence

reflects a fundamental difference in what one-sided flow means across the two contract types. On single-name contracts, the direction of flow predicts the settlement outcome: buckets with YES-directed flow settle YES significantly more often than those with NO-directed flow (logit $z = 12.30$). On broad-based contracts, flow direction and settlement outcome are unrelated (logit $z = -1.74$), because one-sided flow on a Bitcoin or S&P 500 contract reflects the recent trajectory of a continuously moving price that can reverse before settlement. Consistent with the maker profitability decomposition, the VPIN effect on single-name contracts is concentrated on NO-settling contracts ($t = -1.99$), where informed NO flow erodes the optimism-tax profits that fund the maker's operations.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the Kalshi exchange and our data. Section 4 develops our adverse selection measures for the prediction market setting and presents the cross-sectional and panel regression results. Section 5 analyzes maker profitability and decomposes its sources. Section 6 adapts VPIN to prediction markets. Section 7 discusses implications for market design and regulation. Section 8 concludes.

2 Related Literature

The idea that some traders know more than others, and that this asymmetry shapes the cost of trading, is foundational to market microstructure. Copeland and Galai [1983], Glosten and Milgrom [1985], and Kyle [1985] formalize the problem: a market maker who cannot distinguish informed from uninformed order flow must set a spread wide enough to break even against the informed, and the resulting adverse selection cost is borne by all traders. The empirical literature that followed developed increasingly sophisticated tools for detecting information in trade data. Glosten and Harris [1988] decompose spreads into permanent and transitory components; Hasbrouck [1991] measures the information content of trades through their permanent price impact; Easley, Kiefer, O'Hara, and Paperman [1996] develop

the PIN model to estimate the probability that a given trade originates from an informed counterparty. For comprehensive treatments, see O'Hara [1995] and Hasbrouck [2004].

A central prediction of this literature is that adverse selection should be more severe when private information has greater precision. In equity markets, the evidence is consistent: single stocks have wider spreads and higher price impact than index ETFs [Madhavan, 2002], and firms with greater insider ownership or less analyst coverage exhibit elevated adverse selection costs. Our paper tests the analogous prediction in a setting where the cross-sectional variation in information concentration is, if anything, even starker than in equities. Broad-based contracts can be affected by a variety of factors; single event contracts by far fewer. Consequently, we hypothesize that adverse selection should be of greater importance in single-event contracts.

Prediction markets have attracted sustained attention as information aggregation mechanisms. Wolfers and Zitzewitz [2006] survey the evidence that market prices outperform polls and expert surveys, and Arrow [2008] advocates broader adoption. Swanson, Wang, and Wu [2025] show that Kalshi macro contracts rival traditional forecasts for economic variables. The favorite-longshot bias, one of the most robust findings in prediction and betting markets [Wolfers and Zitzewitz, 2004], is directly relevant to our analysis: on Kalshi, it manifests as an "Optimism Tax" in which takers systematically overpay for YES contracts on unlikely outcomes, transferring wealth to makers. Bürgi, Deng, and Whelan [2026] using Kalshi data identify the log-shot bias and heterogeneous beliefs in affecting profit and loss outcomes. Becker [2025] reaches similar conclusions showing that these effects are strongest for contracts on sports and entertainment. Both papers focus on behavioral explanations for wealth transfers. Palumbo [2026] argues that risk allocation is more akin to underwriting than to the intermediation generally found in trading markets. None of these papers measure adverse selection or identify which contracts are most susceptible to informed trading.

A nascent literature has begun to apply microstructure tools to prediction market data, and the early results are suggestive. Tsang and Yang [2026] adapt Kyle's λ to Polymar-

ket election contracts using a log-odds transformation, showing that assassination attempts produce permanent price revisions (information) while debate outcomes produce transitory effects (noise). Hauser and Rittler [2024] apply the Kyle (1985) model to a large prediction market dataset, identifying traders with positive permanent price impact who earn informed-trader profits. Ng, Peng, Tao, and Zhou [2026] find that net order imbalance from large trades predicts returns across multiple platforms.

Our paper explores the extent to which the standard tools of equity market microstructure can be used to assess adverse selection in prediction markets. Facilitating our research are two features of Kalshi's data. First, the taker side field on every trade eliminates the classification error that plagues equity-market adverse selection estimates, allowing clean measurement of signed order flow without the Lee and Ready [1991] algorithm.[5] Second, because we observe both the taker side and the settlement outcome, we can compute the realized profit or loss on every trade for both maker and taker, decompose the sources of maker profitability, and explain why spreads need not widen even when adverse selection is severe.

3 Institutional Setting and Data

3.1 The Kalshi Exchange

Kalshi is a CFTC-regulated designated contract market (DCM) that lists binary and range event contracts. Each binary contract pays \$1 if a specified event occurs and \$0 otherwise; prices are quoted in cents (1–99¢), directly interpretable as market-implied probabilities under risk neutrality.

What distinguishes Kalshi from its predecessors is not the contract design (PredictIt and the Iowa Electronic Markets offered similar binary payoffs) but the market structure. Earlier platforms operated as undifferentiated pools: any participant could post a limit order

[5] See also Easley, López de Prado, and O'Hara [2016] for discussion of alternative methods of trade classification.

or cross the spread, and no formal distinction separated liquidity providers from liquidity consumers.[6] Polymarket, the largest unregulated prediction market, employs a hybrid on-chain order book where sophisticated market makers post limit orders and retail participants typically cross the spread, but the platform does not label trades with maker-taker identifiers. Kalshi, by contrast, operates a continuous double auction with a central limit order book that explicitly tags each trade with the identity of the aggressor. Every trade record discloses the taker side (whether the party who crossed the spread bought YES or NO contracts), and the non-aggressing counterparty is the maker whose resting limit order was executed. This single field eliminates the need for trade-classification algorithms.

Both sides of every Kalshi trade are fully collateralized at execution: the YES buyer deposits the trade price, the NO buyer deposits the complement ($1 minus the trade price), and the exchange holds these funds in segregated accounts at regulated financial institutions. At settlement, Kalshi distributes $1 per contract to the winning side and zero to the losing side. There is no counterparty credit risk: the winner's payout is guaranteed by collateral already posted, not by the loser's ability to pay.[7]

3.2 Market Making Without a Round Trip

The maker in a prediction market is not a neutral intermediary in the sense of Glosten and Milgrom [1985]; she holds a directional position whose payoff depends entirely on the settlement outcome. This distinction from equity market making is fundamental, and it shapes the interpretation of everything that follows.

An equity market maker's business is the round trip. She buys 1,000 shares at $50.00, sells them five minutes later at $50.02, earning 2 cents per share regardless of where the stock trades next month. Her risk is inventory accumulation; the adverse selection problem

[6] PredictIt also capped individual positions at $850 per contract and charged a 10% fee on profits, further blurring any maker-taker distinction.

[7] This contrasts with traditional over-the-counter derivatives, where counterparty risk is managed through margin calls, netting agreements, and central clearing. On Kalshi, the maximum loss is bounded by the initial collateral, and the exchange cannot be called upon to cover defaults.

is that informed traders cause the price to move against her before she can close the position. But the crucial feature is that each trade has a natural offset: the maker's profit comes from buying and selling the same asset.

On Kalshi, there is no round trip.[8] When a maker sells YES at 60 cents, she does not later buy YES to close the position; she holds that exposure until the contract settles. If the event does not occur, she keeps the 60 cents. If it does, she loses 40 cents. She is not earning a spread; she is taking a directional bet at a price she finds attractive. Active position management is not merely inconvenient; it is economically untenable. In equity markets, closing a position is capital-*releasing*: the maker sells the shares she bought, frees the capital, and redeploys it. On Kalshi, full collateralization makes closing capital-*intensive*: "closing" a YES sale by buying YES on the same contract requires posting *additional* collateral equal to the YES price, increasing the capital commitment rather than freeing it.[9] Prediction market contracts have no close substitutes for hedging (a maker short YES on a Tesla delivery contract cannot hedge with Tesla stock, because the payoff depends on a discrete threshold, not the continuous stock price), and contracts have finite, often short, lives that limit the benefit of dynamic management relative to simply holding to settlement.

Losing on individual trades is therefore structural, not pathological. A maker who sells YES at 80 cents on a contract that the market correctly prices at an 80% probability fully expects to lose this trade four times out of five, losing 20 cents each time. The one time in five she wins, she wins 80 cents; in expectation, the trade breaks even. The math of binary settlement guarantees losers. Her edge, if she has one, comes from the *other* trades on the same contract: takers who buy YES at 20 cents on a 20% event, or who buy NO when the market is overpricing the outcome. If enough takers are on the wrong side, the maker's wins

[8] A market maker could in principle offset positions through later trades; for example, it could sell YES early in the life of a contract and later buy YES or sell NO as order flow arrives. In practice this mechanism is limited. Because contracts are fully collateralized, offsetting trades typically require posting additional collateral rather than releasing capital. Moreover, contracts are idiosyncratic and short-lived, leaving few hedging instruments or substitutes. As a result, inventory reduction depends largely on endogenous order flow rather than active position management, and makers often carry substantial exposure until settlement.

[9] Kalshi does allow position netting in some cases, but the maker must still cross the spread as a taker to execute the offset, paying the bid-ask spread rather than earning it.

on those trades cover her expected losses on the correctly priced ones.

When a maker loses money on a contract as a whole, one of two things went wrong: the flow was too one-directional for her to accumulate offsetting trades (the informed-trading channel), or the price was systematically wrong about the outcome (the mispricing channel). In practice, these blend together, because informed traders pushing flow in one direction are precisely what prevents the maker from building the offsetting positions she needs.

What, then, does "adverse selection" mean in this setting? In equity markets, it is about *price*: the informed trader causes the price to move against the market maker before she can close her position. In prediction markets, it is about *outcome*: the informed trader is on the correct side of the settlement, and the maker cannot undo the trade. There is no position to unwind, no spread to earn on the round trip, no inventory to dynamically manage. The measures we develop in Section 4 (Kyle's λ and the Glosten-Harris decomposition) detect the same phenomenon (informed order flow moving prices), but the economic mechanism through which it harms the maker is fundamentally different. The frequency-magnitude decomposition in Section 5 is designed to capture this difference. The VPIN toxicity analysis in Section 6 refines this exposure for the market maker.

This institutional distinction also creates a rare empirical opportunity. Because we see both the taker side and the settlement outcome for every trade, we can compute the realized profit or loss for both maker and taker on every transaction without assumptions about inventory management, position offsets, or market-maker identity.

3.3 Data

We access all trade, market, and candlestick data via Kalshi's REST API. The sample spans the exchange's full operating history from July 2021 through March 11, 2026, and contains over 7 million traded markets organized into nearly 4 million event series. Figure 1 illustrates the exchange's growth: monthly dollar volume rose from under $1 million in mid-2021 to over $2.6 billion by November 2025. Across all markets, we observe 258.7 million trades on 7.4

million contracts. For midpoint construction, we use hourly candlestick bid-ask snapshots. These valid quotes (where bid > 0, ask < 100, and ask $>$ bid) are available for approximately 79% of trade-candlestick matches in the analysis sample. Trades that fall outside candlestick coverage are excluded from all spread-based calculations.

3.4 Contract Classification

The classification that organizes our analysis is simple: we separate contracts by how concentrated private information about their outcomes is likely to be.

Single-name contracts reference outcomes specific to a single company or individual: company earnings-call keyword mentions, Tesla delivery targets, Netflix subscriber milestones, CEO departures, IPO outcomes, Super Bowl celebrity attendance, and similar events. Private information about these outcomes is concentrated among a small number of insiders: corporate executives, employees, family members, and their networks.

Broad-based contracts reference macroeconomic indicators, market indices, and other outcomes where information is dispersed across many agents: S&P 500 ranges, Nasdaq 100, Federal Reserve rate decisions, CPI, GDP, unemployment, Bitcoin and Ethereum ranges, Treasury yields, and foreign exchange rates. Private information here is likely to be on factors contributing to the outcome.

This classification parallels the single-stock versus index-ETF comparison in equity markets. A corporate insider who knows Tesla's quarterly delivery numbers has a decisive advantage over market makers quoting Tesla contracts; no individual possesses comparable private information about whether the S&P 500 will close above 5,000 tomorrow. Appendix A sets out the classification rules and the specific contract series assigned to each group. The remaining markets (primarily politics, sports, entertainment, and weather) do not map cleanly onto either category and are excluded from the main analysis.

Table 1 reports summary statistics for the 478,167 single-name and broad-based markets with at least one recorded trade, which constitute the analysis sample for the rest

of the paper. Single-name contracts account for 15,516 of these markets and 3.0 million trades; broad-based contracts account for 462,651 markets and 38.6 million trades. Takers on single-name contracts buy YES 60.2% of the time, compared with 53.2% on broad-based, consistent with the behavioral YES bias we document in Section 5. Settlement rates differ markedly: only 30.9% of single-name contracts settle YES versus 38.4% of broad-based. The gap between the YES-buying rate and the YES-settlement rate (29.3 percentage points on single-name versus 14.8 on broad-based) foreshadows the frequency edge that drives maker profitability.

[Insert Table 1 about here]

3.5 Three Examples: What Happens to Makers?

Two contracts from the Super Bowl LX celebrity attendance series illustrate the adverse selection problem, and a third shows what happens when it overwhelms the maker.

The Jeff Bezos contract (KXSBGUESTS-26-JBEZ) opened at 62 cents and briefly spiked to 95 cents on uninformed YES speculation. The *Wall Street Journal* documented what happened next: Evan Whitesell, Bezos's stepson and a member of Sigma Alpha Epsilon at the University of Miami, had discussed his family's Super Bowl plans at a fraternity event, revealing that Bezos would *not* attend. Fraternity brothers and their networks bet NO, driving the YES price from roughly 70 to 30 cents over several days. The contract settled NO; the informed traders were correct. Yet despite facing genuine insider trading, makers earned $29,274 in total (an average of 3.3 cents per contract) because uninformed YES-biased takers (518,000 contracts at an average of 27 cents) more than offset the losses imposed by informed NO flow (370,000 contracts).

The Mark Wahlberg contract (KXSBGUESTS-26-MWAH) tells the complementary story. A rumor originating from a Clemson University fraternity generated 24 million contracts traded ($8.8 million in notional volume) on the question of whether Wahlberg would attend. The rumor was false. Takers were 68.6% YES-biased, paying an average of 37 cents per YES

contract for an outcome that never materialized. Makers earned $1.18 million. No insiders, no private information—just behavioral demand generating massive transfers from takers to makers.

The contrast is sharp. On the Bezos contract, insider trading was real; on the Wahlberg contract, the information was illusory. In both cases, makers profited. Figure 2 plots the price paths for both contracts, illustrating the insider-driven price *decline* on Bezos and the rumor-driven volume surge on Wahlberg. Figure 3 shows the cumulative maker P&L on the Bezos contract: it fluctuates as informed NO traders impose losses, but the cumulative effect of uninformed YES-biased takers dominates, ending at $29,274 at settlement. The question our paper addresses is whether this pattern (maker profitability despite adverse selection) holds systematically across the full Kalshi universe.

Not all single-name contracts end well for makers. The Netflix Q4 2025 earnings-mention contract "Will Netflix say 'Warner Bros'?" (KXEARNINGSMENTIONNFLX-26JUN30-WARN) illustrates the opposite extreme. During the live earnings interview (4:45–5:30 PM ET on January 20, 2026), traders watching the broadcast in real time overwhelmed makers with NO orders: 210,000 contracts of NO-side volume versus 71,000 YES-side, at an average yes price of 61 cents. As each minute passed without the phrase being uttered, the NO takers' information advantage grew. Makers lost $85,407 during the 45-minute interview alone and $75,754 over the contract's lifetime—the single worst contract for makers in our entire sample. Figure 4 displays the cumulative maker P&L alongside the YES price, with vertical lines marking the earnings release and the interview window.

The "Warner Bros" contract also reveals that adverse selection in single-name contracts operates through multiple channels beyond classical informed trading. *Real-time observation* of a public broadcast is one: anyone watching the livestream has a minutes-long advantage over makers whose quotes are stale. Another is *contract-design risk*: a Netflix executive said "Warner Brothers" during the interview, but Kalshi's rules require the exact phrase "Warner Bros," and the contract resolved NO, rewarding sophisticated traders who understood the

literal resolution criteria over novices who assumed the spoken reference was sufficient. Both channels are concentrated in single-name contracts and largely absent from broad-based ones.

[Insert Figures 2, 3, and 4 about here]

4 Adverse Selection in Prediction Markets

4.1 Metrics

Measuring adverse selection in prediction markets requires adapting the standard microstructure toolkit. We employ three measures: the effective spread, which captures the cost of immediacy paid by takers; Kyle's λ, which estimates the permanent price impact of signed order flow; and the Glosten and Harris [1988] decomposition, which separates the price impact into permanent (information) and transitory (inventory) components. Each requires some modification for the prediction market setting. Prediction market prices are bounded between 0 and 100 cents, creating heteroskedasticity that distorts linear regressions in levels. Quote data are available only at hourly frequency, limiting the precision of midpoint-based measures. And the absence of a round trip means that the economic interpretation of each measure differs from its equity-market counterpart.

One standard approach, however, is unavailable to us entirely. In equity markets, the Huang and Stoll [1996] decomposition splits the effective spread into a realized spread (the maker's short-term profit) and an adverse selection component (the post-trade midpoint movement in the direction of the taker's trade), estimated at horizons of 5–30 minutes after each trade. Our hourly candlestick data are too coarse for this purpose: 73% of trades have an identical candlestick midpoint at the 5-minute horizon (because the next observation falls within the same hourly candle), and 52% are identical at 30 minutes. The resulting adverse selection component is mechanically zero for the majority of observations. We therefore rely on Kyle's λ and the Glosten-Harris permanent component, both estimated from trade prices and signed order flow with no dependence on forward quote midpoints, as our primary mea-

sures of informed trading. The effective spread remains useful as a measure of trading costs, even though its decomposition into realized and adverse selection components is infeasible.

4.1.1 Effective Spread

The effective half-spread measures the cost of immediacy: how much the taker pays above (or below) the prevailing midpoint. For each trade i on contract j with taker side $q_i \in \{+1, -1\}$ (YES $= +1$, NO $= -1$) and midpoint M_i computed from the enclosing hourly candlestick bid-ask data:

$$\mathrm{ES}_i = q_i \cdot (P_i - M_i)$$

where P_i is the `yes_price` of the trade. We compute midpoints from the closing bid and ask of the enclosing hourly candlestick; trades that fall outside candlestick coverage (approximately 21% of the sample) are excluded from all spread-based calculations.

4.1.2 Kyle's λ in Log-Odds Space

Kyle's λ captures the impact of signed order flow on price changes—a higher λ is consistent with greater informed trading. In prediction markets, however, the bounded price space ($[0, 100]$ cents) creates a problem: a 1-cent price change near 50 cents has very different information content than a 1-cent change near 5 cents. Following Tsang and Yang [2026], we transform prices to the log-odds (logit) scale:

$$\ell_{j,h} = \log\left(\frac{p_{j,h}/100}{1 - p_{j,h}/100}\right)$$

which maps $p_{j,h} \in (0, 100)$ to $\ell_{j,h} \in (-\infty, +\infty)$, removing the boundary effects. On this scale, equal absolute changes in log-odds correspond to equal proportional changes in the odds ratio regardless of the price level.

For each contract j, we aggregate trades into 5-minute intervals indexed by h and compute signed order flow as $\mathrm{SOF}_{j,h} = \sum_{i \in h} q_i \cdot c_i$, where $q_i \in \{+1, -1\}$ is the taker sign and c_i is the

number of contracts. Kyle's λ_j is the slope from the contract-level time-series regression:

$$\Delta\ell_{j,h} = \alpha_j + \lambda_j \cdot \mathrm{SOF}_{j,h} + \varepsilon_{j,h}$$

where $\Delta\ell_{j,h} = \ell_{j,h}^{\text{close}} - \ell_{j,h}^{\text{open}}$ is the within-interval change in log-odds. A higher λ_j means that a given quantity of net order flow moves prices more, the signature of informed trading. We require a minimum of 50 five-minute intervals with trades per contract; contracts with insufficient trading activity are excluded from all Kyle's λ and Glosten-Harris estimates.

4.1.3 Glosten-Harris Decomposition

The Glosten-Harris decomposition separates price impact into a permanent component (the portion that reflects new information and is not reversed) and a transitory component (the portion that captures inventory or behavioral effects). Following Glosten and Harris [1988], for each contract j:

$$\Delta p_{j,h} = \gamma_{0,j} + \gamma_{1,j} \cdot \mathrm{SOF}_{j,h} + \gamma_{2,j} \cdot \Delta\mathrm{SOF}_{j,h} + \varepsilon_{j,h}$$

where $\Delta\mathrm{SOF}_{j,h} = \mathrm{SOF}_{j,h} - \mathrm{SOF}_{j,h-1}$. The permanent component $\gamma_{1,j}$ is the coefficient on the level of order flow; the transitory component $\gamma_{2,j}$ is the coefficient on its first difference. We estimate the regression in both raw-price and log-odds space.

The Glosten-Harris specification nests the Kyle regression (it simply adds ΔSOF as a second regressor), so $\gamma_{1,j}$ and λ_j are mechanically related. If SOF and ΔSOF are uncorrelated, $\gamma_{1,j} = \lambda_j$ exactly; in our data, the cross-contract correlation between the two estimates is 0.91 in log-odds space, and the mean R^2 improvement from adding ΔSOF is just 0.0002. We report the Glosten-Harris results not as an independent test of informed trading but for what the transitory component reveals.

The answer is striking. The transitory component $\gamma_{2,j}$ is economically negligible (mean ≈ 0) and statistically significant in only 6% of contracts. Virtually all of the price impact

from order flow in prediction markets is *permanent*, consistent with information incorporation rather than short-lived inventory or liquidity effects. In equity markets, by contrast, the transitory component is typically large, reflecting market makers' active inventory management and mean-reverting quote adjustments [Huang and Stoll, 1996]. The near-absence of a transitory component here is exactly what the institutional setting predicts: prediction market makers hold directional positions to settlement rather than managing inventory dynamically, so there is no inventory effect to reverse.

4.2 Results

4.2.1 Cross-Sectional Regressions

We begin with a simple question: do contracts where private information is likely concentrated among a small number of insiders exhibit more informed trading than contracts where information is widely dispersed? Our main specification regresses contract-level adverse selection measures on an indicator for single-name contracts:

$$\text{AS}_j = \alpha + \beta_1 \cdot \text{SINGLE_NAME}_j + \beta_2' X_j + \varepsilon_j$$

where X_j includes log volume, price level, price level squared, and implied volatility ($\sqrt{p(1-p)}$). Standard errors are heteroskedasticity-robust (HC1). The hypothesis is straightforward: if information concentration matters, $\beta_1 > 0$.

[Insert Table 2 about here]

Table 2 reports the results.[10] Kyle's λ is significantly higher for single-name contracts in both raw-price ($\beta = 0.0010$, $t = 9.20$) and log-odds ($\beta = 0.0001$, $t = 10.88$) specifications

[10]The observation counts differ across tables because each measure imposes different data requirements. Table 1 reports the 478,167 single-name and broad-based markets with at least one trade (15,516 + 462,651). Kyle's λ and the Glosten-Harris decomposition require at least 50 five-minute trading intervals per contract, reducing the sample substantially; most low-activity contracts have too few intervals for reliable estimation. The effective spread requires only a valid candlestick midpoint, so column (5) retains more contracts. Table 5 further restricts to settled contracts, since maker profit cannot be computed without a settlement outcome.

(columns 1 and 2). The Glosten-Harris permanent component tells the same story ($t = 7.46$ and $t = 9.58$ in columns 3 and 4). As noted in Section 4, these estimates are not independent of Kyle's λ ($\rho = 0.91$), but they confirm that virtually all price impact is permanent. Column (5) shows that effective spreads are significantly wider for single-name contracts ($\beta = 0.29$, $t = 3.17$).

The results are exactly what microstructure models would predict. Contracts referencing outcomes known to a handful of insiders exhibit more informed trading, that information is incorporated permanently into prices, and spreads are wider. Our main analysis focuses on single-name and broad-based contracts because the cross-sectional variation in information concentration is sharpest at these two extremes. But Kalshi lists millions of additional contracts across politics, sports, entertainment, weather, and other categories. In Appendix B, we extend the Kyle λ analysis to the full universe of non-parlay Kalshi contracts (129,731 contracts across seven categories) and find that the single-name/broad-based dichotomy does in fact capture the two poles of informed trading: single-name contracts exhibit the highest λ of any category, and broad-based contracts among the lowest, with politics, sports, entertainment, and other categories falling in between.[11]

But in prediction markets, wider spreads do not protect the maker the way they do in equities. As discussed in Section 3, the prediction market maker holds a directional position to settlement; her profit depends on the settlement outcome, not on the bid-ask spread. The next result shows that spreads alone do not explain maker profitability, because the prediction market maker's profit comes from the gap between the trade price and the settlement value, not from a round-trip spread.

[11]The one apparent exception is weather contracts, whose overall λ exceeds that of single-name contracts. A directional decomposition that separates the price impact of order flow on the correct side of the settlement outcome from flow on the wrong side reveals that this anomaly is driven entirely by uninformed flow that moves prices but is systematically wrong about the terminal value. On the dimension that matters for informed trading, single-name contracts dominate every other category.

4.2.2 Robustness

Table 3 subjects the main finding to several alternative specifications, all using the log-odds Kyle's λ. The single-name effect survives each one. It is highly significant without any controls ($t = 14.39$, column 1) and within the financial-contracts-only subsample that excludes politics, sports, and entertainment ($t = 9.75$, column 2). Replacing the binary single-name indicator with a continuous measure of insider concentration yields an equally significant coefficient ($t = 11.92$, column 3), confirming that the result is not an artifact of the classification boundary.

A separate concern is midpoint quality. Our effective spread measure relies on the closing bid and ask of the enclosing hourly candlestick, so a trade at minute :05 of the hour is matched to a quote observed 55 minutes later. If this staleness biases our spread estimates, it could mask a true difference between single-name and broad-based contracts. To assess this, we restrict the sample to trades occurring within τ minutes of the end-of-period quote and re-estimate the cross-sectional spread regression.

[Insert Tables 4 and 3 about here]

Table 4 reports the results. The single-name effective spread premium is positive across all windows, increasing from $\beta = 0.29$ ($t = 1.60$) at $\tau = 5$ minutes to $\beta = 0.17$ ($t = 2.27$) for the full sample. The pattern is consistent with wider spreads on single-name contracts regardless of midpoint staleness. As the window narrows and the sample shrinks, statistical significance declines but the sign and magnitude remain stable.

4.2.3 The Limits of Spread Compensation

Wider spreads on single-name contracts are the conventional prediction of adverse selection theory: informed order flow raises the cost of immediacy, and liquidity providers respond by pricing less aggressively. In equity markets, this response also compensates the maker. A wider spread means the maker buys at a lower bid and sells at a higher ask, earning a

larger markup on each uninformed trade. Because the loss from an informed trade is typically comparable in magnitude to the spread (a few cents of adverse price movement on a stock that moves in small increments), the markup on uninformed trades can, in equilibrium, cover the losses on informed ones. The canonical Glosten and Milgrom [1985] equilibrium is exactly this balance.

In prediction markets, wider spreads provide the same benefit at the trade level: a maker who sells YES at 57 cents instead of 55 has a better entry price, earning 2 additional cents if the contract settles NO and losing 2 fewer cents if it settles YES. The half-spread helps by exactly 2 cents per trade, just as it does in equities. The difference is in the loss side. An equity market maker who sells at \$55 and faces an informed buyer typically loses a few dollars when the true value turns out to be \$58 or \$60, a loss that is a small multiple of the spread. A prediction market maker who sells YES at 57 cents and faces an informed buyer who knows the event will occur loses $100 - 57 = 43$ cents—the full distance to settlement. Binary settlement creates an asymmetry that has no analog in equity markets: the maker's gain per uninformed trade is bounded by the spread, but her loss per informed trade is bounded by the distance between the trade price and the terminal value of 0 or 100.

To see how demanding this asymmetry is, consider a setting in which a fraction α of takers are informed and the remaining $1-\alpha$ are uninformed. For a contract priced at P with half-spread s, the maker breaks even only if

$$s \geq \alpha \cdot (100 - P).$$

At $P = 50$ with $\alpha = 0.20$, the required half-spread is 10 cents, a 20-cent full spread. Empirical effective half-spreads in our data average approximately 2 cents, sufficient to offset at most $s/(100-P) = 2/50 = 4\%$ informed takers. Even this overstates the spread's capacity, because the empirical spread must also cover inventory carrying costs and order processing, leaving less margin for adverse selection losses. At extreme prices the constraint tightens further:

the 0–100 cent pricing grid limits the maximum feasible half-spread to $\min(P, 100 - P)$, so at $P = 90$ the widest possible half-spread is 10 cents, which can absorb at most $10/90 \approx 11\%$ informed takers regardless of how aggressively the maker prices. In practice, some makers do accumulate offsetting positions through later order flow, reducing the effective loss per informed trade below this worst case. But the scale of the mismatch (observed spreads would need to be five times wider to break even at even moderate levels of informed trading) indicates that spread revenue alone cannot sustain market making on contracts where informed trading is a significant component of order flow.

Spreads on Kalshi therefore reflect adverse selection risk (wider on single-name contracts, where informed trading is more severe), but the spread is not the mechanism through which makers are compensated for that risk. If the spread cannot do the job, what does?

5 Maker Profitability

Single-name contracts have wider spreads and higher price impact, confirming the presence of adverse selection. But the analysis above shows that wider spreads cannot protect prediction market makers the way they protect equity dealers. How, then, do makers profit on the very contracts where informed trading is most severe? In this section, we answer the question by computing maker profits directly from settlement outcomes.

5.1 Computing Maker Profits

Because we observe the taker side and settlement outcome for every trade, we can compute the exact maker profit on each transaction:

$$\pi_i^{\text{maker}} = \begin{cases} (P_i - V_j) \cdot c_i & \text{if taker_side} = \text{YES (maker is short YES)} \\ (V_j - P_i) \cdot c_i & \text{if taker_side} = \text{NO (maker is long YES)} \end{cases}$$

where P_i is the yes_price (in cents), $V_j \in \{0, 100\}$ is the settlement value, and c_i is the number of contracts. No assumptions about market-maker identity, inventory, or position limits are required. The computation simply asks: for every trade, what did the non-aggressing side earn at settlement?[12]

This measure assumes that the maker holds each position to settlement rather than offsetting it through subsequent trading. As discussed in Section 3, the assumption is well-grounded: full collateralization makes offsetting trades capital-intensive rather than capital-releasing, contracts have no close substitutes for hedging, and short contract lives limit the benefit of dynamic management. Holding to settlement is not merely a convenient simplification; it is a reasonable description of how most makers operate in this market. Our per-trade cents measure also preserves the zero-sum accounting that is fundamental to a two-sided market: every cent the maker earns, the taker loses, and vice versa.[13]

Even if some makers offset positions before settlement, the measure remains informative for cross-sectional comparisons. Consider a maker who sells YES at 60¢ on a contract that settles NO. Our computation credits a 60-cent profit regardless of whether she held the position or sold it earlier at 40¢ (realizing only 20 cents). These measurement errors attenuate toward zero across our large sample, and the cross-sectional comparison between single-name and broad-based contracts is valid as long as offsetting behavior does not differ systematically between the two groups.[14]

[12]Our analysis ignores trading fees imposed by Kalshi. Kalshi charges taker fees on all contracts, ranging from $0.07 to 1.75$ per contract depending on trade size and contract type. Starting in April 2025, makers on some contract types also face fees ranging from $0.02 to 0.44$. Given the complexity and time-varying nature of the fee schedule over our sample period, we do not incorporate fees in our calculations. To the extent that fees reduce maker profits, our estimates overstate true profitability for both contract types; the cross-sectional comparison between single-name and broad-based contracts is unaffected unless fee structures differ systematically across the two groups.

[13]Bürgi, Deng, and Whelan [2026] report that both makers and takers earn negative average *percentage* returns on Kalshi. This does not violate the zero-sum constraint; it reflects the favorite-longshot bias in an equal-weighted average across price levels. A longshot at 5¢ that loses costs −100%; a favorite at 95¢ that wins returns only +5.3%. Because two-thirds of Kalshi contracts trade below 10¢ or above 90¢, the massive percentage losses on longshots drag the average return negative for both sides. Our cents-per-trade decomposition avoids this artifact and makes the cross-subsidy from uninformed to informed trading directly visible.

[14]Active inventory management is more likely on broad-based contracts, where correlated contracts (e.g., adjacent S&P 500 range brackets) could serve as partial hedges. If anything, this would cause our measure

5.2 Maker Profits by Contract Type

Table 5 reports maker profitability for the 475,646 settled single-name and broad-based contracts with trading activity, and the headline result is striking.[15]

[Insert Table 5 about here]

Makers earn far more on single-name contracts—the very contracts where adverse selection is most severe. Per-trade profits average 3.94 cents on single-name versus just 0.74 cents on broad-based, over five times as much. The maker win rate is 56.6% on single-name but only 50.7% on broad-based. In dollar terms, single-name makers earn $305.89 per contract compared to $69.89 for broad-based, and 66.5% of single-name contracts are profitable for the maker versus only 49.4% of broad-based. This is difficult to reconcile with the standard view that adverse selection erodes maker profitability.

5.3 The Frequency-Magnitude Decomposition

Why do makers earn more where informed traders are most active? We decompose expected maker profits into a frequency component and a magnitude component:

$$\begin{aligned} E[\pi^{\text{maker}}] &= \Pr(\text{win}) \cdot E[\pi|\text{win}] - \Pr(\text{loss}) \cdot E[|\pi| \,|\, \text{loss}] \\ &= \underbrace{(\text{WR} - 0.5)(W + L)}_{\text{frequency edge}} + \underbrace{0.5(W - L)}_{\text{magnitude edge}} \end{aligned}$$

where WR is the win rate, W is the average win, and L is the average loss magnitude.

[Insert Table 6 about here]

The decomposition reveals that maker profits on single-name contracts are overwhelmingly driven by the frequency edge (Figure 5). Of the 3.94-cent average per-trade profit, 3.85

to *overstate* broad-based maker profits relative to the true hold-to-settlement benchmark, biasing against our finding that single-name makers earn more per trade.

[15]Maker profit requires a settlement outcome, so unsettled contracts are excluded. The 475,646 contracts here (14,988 single-name + 460,658 broad-based) represent the subset of Table 1's universe that has both settled and recorded at least one trade.

cents (97.7%) comes from the frequency edge: makers win 56.6% of trades because takers are disproportionately YES-biased (60.3%) on contracts where only 30.9% settle YES. Win and loss magnitudes are nearly identical (29.24 versus 29.06 cents); the magnitude edge contributes a negligible 0.09 cents. The frequency edge is doing virtually all of the work.

Broad-based contracts show a qualitatively similar pattern but at dramatically smaller scale. The frequency edge (0.41 cents) accounts for 55.6% of the 0.74-cent per-trade profit, with the magnitude edge contributing the remaining 0.33 cents. Broad-based makers win a bare majority (50.7%) of trades and earn modestly more when they win (31.80 cents) than they lose when they lose (31.14 cents). Both edges are an order of magnitude smaller than on single-name contracts, reflecting the absence of a large behavioral surplus to harvest.

[Insert Figure 5 about here]

5.4 Why Is the Win Rate Higher in Single-Name?

The arithmetic of the maker's win rate is illuminating. The maker wins in exactly two cells of the taker-side × settlement table: (A) the taker buys YES and the contract settles NO, or (B) the taker buys NO and the contract settles YES. Under independence between taker direction and settlement:

$$\widehat{\text{WR}} = \Pr(\text{taker=YES}) \times \Pr(\text{settles NO}) + \Pr(\text{taker=NO}) \times \Pr(\text{settles YES})$$

For single-name contracts, the predicted win rate under independence is $0.603 \times 0.691 + 0.397 \times 0.309 = 53.9\%$. The actual win rate is 56.6%, a gap of 2.7 percentage points that reveals something important. Conditional on the marginal distributions, takers are *anti-informed*: they buy YES on contracts that settle NO, and sell YES on contracts that settle YES, even more than the marginal YES-bias alone would predict. For broad-based contracts, the predicted and actual win rates are virtually identical: $0.532 \times 0.616 + 0.468 \times 0.384 = 50.7\%$, compared to an actual rate of 50.7%. Broad-based takers are, on average, well-

calibrated (neither informed nor anti-informed), which is why the behavioral surplus available to broad-based makers is so much smaller.

Two underlying forces drive the higher single-name maker profitability: takers have a strong YES-bias (60.3%) on contracts that settle YES only 30.9% of the time, and the bulk of single-name takers are behaviorally motivated. The *average* single-name taker is anti-informed even though a *subset* (the insiders) are informed. The fraternity brothers who knew Bezos would not attend the Super Bowl were genuine informed traders; the tens of thousands of YES takers who thought he would attend were not.

5.5 The Longshot Bias as Cross-Subsidy

Where, exactly, does the behavioral surplus come from? The maker P&L by taker price bucket (Table 7) provides a detailed answer.

[Insert Table 7 about here]

On single-name contracts, maker P&L is positive across nearly every price bucket, peaking at +9.08 cents for takers paying 31–40 cents. Even at prices near the midpoint, where information content should be highest, makers earn substantial positive returns (8.86 cents at 41–50 cents, 5.31 cents at 51–60 cents). Only in the 71–80 cent range does maker P&L turn slightly negative (−0.51 cents). The pattern suggests that adverse selection losses from informed traders are absorbed and more than recovered from uninformed takers at similar price points.

Broad-based contracts show a starker divide. Maker P&L is positive at low taker prices (1–60 cents) but turns negative above 60 cents, where takers are buying near-certain favorites. At these price levels there is little behavioral surplus to harvest, and even mildly informed takers can impose losses.

We can quantify the dependence on the longshot bias by splitting trades at a 25-cent taker price threshold. On single-name contracts, 59.7% of total maker P&L comes from longshot

takers (those paying ≤ 25 cents). On broad-based contracts, the dependence is even more extreme: *all* of maker P&L comes from longshot trades (104.2%), because non-longshot trades actually generate net losses (-4.2%). Both contract types depend on behavioral demand from takers who overpay for unlikely outcomes, but the dependence is particularly acute for broad-based contracts. If the longshot bias were to diminish (through better-calibrated takers, algorithmic pricing tools, or regulatory restrictions on extreme-probability contracts), broad-based maker profits would vanish entirely, and single-name maker profits would be substantially reduced.

5.6 Contract-Level Profitability Distribution

Averages, of course, mask considerable heterogeneity. Figure 6 displays the distribution of total maker P&L across individual contracts. Among single-name contracts, 66.5% are profitable and 33.5% generate losses; among broad-based contracts, the split is essentially even at 49.4% profitable versus 50.4% unprofitable. Both distributions are right-skewed (mean maker P&L is \$306 on single-name (median \$4) and \$70 on broad-based (median \$$-$0)), reflecting a small number of high-volume contracts that generate outsized profits.

The contrast is stark: the typical broad-based contract is essentially a coin flip for the maker, while the typical single-name contract generates a small profit. Makers losing money on half of broad-based contracts underscores that adverse selection risk is not theoretical; it materializes at the contract level with high frequency. The maker's edge is a *portfolio* phenomenon: across many contracts, the behavioral surplus from uninformed takers more than offsets the losses on contracts where informed traders dominate. The Netflix "Warner Bros" contract from Section 3 (where makers lost \$75,754) is the tail risk; the Wahlberg contract (where makers earned \$1.18 million from a false rumor) is the other side of the same distribution.

[Insert Figure 6 about here]

5.7 Where Informed Trading Hurts: Settlement Outcome and the Optimism Tax

Where in the maker's book do informed traders inflict damage? Because we observe both the taker side and the settlement outcome for every trade, we can decompose total maker P&L into four cells of the taker-side × settlement table. Table 8 reports the results.

[Insert Table 8 about here]

The maker's entire profit comes from contracts that settle NO. On single-name contracts, NO-settling contracts generate $6.25 million in maker profits, and 78% of these contracts are profitable for the maker. YES-settling contracts lose $1.66 million, with only 41% profitable. Broad-based contracts display the same asymmetry: +$64.5 million on NO-settling versus −$32.3 million on YES-settling. Takers buy YES 68.7% of the time on single-name contracts that settle NO, so the maker keeps the premium when those contracts expire worthless.

Within NO-settling contracts, the trade-level decomposition reveals a tug-of-war between two populations. YES takers, the uninformed optimists whose bets fail at settlement, generate $18.2 million in single-name maker profits, fully 397% of net income. NO takers on the same contracts, the informed traders who correctly bet against the event, impose losses of $12.0 million, or −261% of net. The maker's profit on NO-settling contracts is the residual after these two forces offset. The optimism tax exceeds the informed-trading cost by a factor of 1.52× on single-name and 1.28× on broad-based.

The maker's strategy works only so long as the optimism tax exceeds the cost of trading against informed counterparties on contracts that settle NO. Losses on YES-settling contracts are structural: some events do occur, and the maker is on the wrong side when they do. The real risk is that informed traders erode the edge on NO-settling contracts. To see this directly, we sort NO-settling contracts by the fraction of taker volume that is NO-side and examine maker outcomes across quintiles. On single-name contracts, the lowest quintile of NO-taker flow (mean NO fraction = 0%, where every taker bought YES) has a maker loss rate of 0%.

The highest quintile (mean NO fraction = 80%, dominated by correct NO bettors) has a loss rate of 77%. On broad-based contracts, the gradient is equally sharp: 0% to 100% (logit $z = 294$, $p < 0.001$).

So why do single-name makers earn more despite facing more informed trading? Because the optimism tax on single-name contracts is large enough to absorb it: the ratio of optimism-tax revenue to informed-trading cost is 1.52× on single-name versus 1.28× on broad-based. Whether that margin holds on any given contract at any given moment is the question we take up next.

6 Flow Toxicity: VPIN in Prediction Markets

Can the toxicity of order flow be detected in real time? The Volume-Synchronized Probability of Informed Trading (VPIN), developed by Easley, López de Prado, and O'Hara [2012], measures the degree to which order flow is one-sided within equal-sized volume buckets and has been shown to identify elevated toxicity in the hours preceding the May 2010 Flash Crash [Easley, López de Prado, and O'Hara, 2011]. Adapting VPIN to prediction markets requires several modifications, but the resulting metric successfully identifies when informed trading is eroding the maker's profits on single-name contracts.

6.1 Adapting VPIN to Prediction Markets

VPIN partitions trading activity into n volume buckets of equal size V, classifies each trade as buyer- or seller-initiated, and computes:

$$\text{VPIN} = \frac{1}{n} \sum_{\tau=1}^{n} \frac{|V_\tau^B - V_\tau^S|}{V_\tau}$$

where V_τ^B and V_τ^S are the buy and sell volumes in bucket τ. VPIN ranges from 0 (perfectly balanced flow) to 1 (completely one-sided flow).

Three features of Kalshi's market structure require modification of the standard equity VPIN. First, trade classification is exact: the taker side field directly identifies whether the aggressor bought YES or NO, eliminating the classification error inherent in the bulk volume classification or Lee-Ready approaches used in equity markets. Second, volume must be measured in dollars rather than contracts, because a trade of 1,000 contracts at 90¢ deploys nine times the capital of the same trade at 10¢; using contract counts would treat these as equivalent.[16] Third, bucket sizing must accommodate extreme trade-size skew. The median Kalshi trade is 21 contracts, but the 99th percentile is 1,500 and the maximum exceeds 500,000 (skewness = 153). The standard VPIN bucket size of average daily volume divided by 50 produces buckets so small that individual whale trades fill entire buckets, mechanically pushing VPIN toward 1.0. With the standard parameterization, mean VPIN exceeds 0.96 across all contract types, too compressed to discriminate between informed and uninformed flow. We address this by using large, fixed dollar-amount buckets ($500 in our baseline specification) and requiring each contract to generate at least five full buckets for inclusion.[17]

A fourth issue is specific to binary settlement and has no equity-market analog. As a contract approaches its settlement date and the outcome becomes clear, the price converges toward 100 (for YES-settling contracts) or 0 (for NO-settling contracts). This convergence mechanically generates one-sided flow: nearly all late trades push in the direction of the terminal value. The resulting VPIN inflation is not a signal of informed trading; it is a consequence of price convergence that would occur even if every trader were uninformed. In our data, the YES-settlement rate among single-name contracts in the highest VPIN quintile is 95%, compared to 24% in the lowest quintile. The apparent relationship between VPIN and maker losses in the pooled data is substantially driven by this settlement-outcome

[16] We also examined raw contract counts and max-loss exposure ($c_i \times \min(P_i, 100 - P_i)/100$). Dollar notional produces the strongest and most monotonic relationship with maker outcomes. Results for all three volume definitions are available from the authors.

[17] We sweep bucket sizes from $100 to $50,000. The results are qualitatively robust. At $500, the sample includes 4,312 single-name and 77,329 broad-based contracts. We report the $500 specification as the best balance of sample size and discriminating power.

composition effect.[18]

We address this by restricting the analysis to trades executed at interior prices, defined as a volume-weighted average yes_price between 30 and 70 cents within the bucket. At these prices the maker's exposure is economically meaningful, because neither side of the trade is nearly certain to win, and the mechanical link between VPIN and settlement outcome is broken.[19] A contract trading at 50 cents is roughly equally likely to settle YES or NO regardless of whether recent flow has been one-sided, so any relationship between VPIN and maker losses at interior prices reflects genuine flow toxicity rather than price convergence.

6.2 Trailing VPIN as a Real-Time Predictor

We construct VPIN as a *trailing* indicator that a maker could monitor in real time. For each contract, we partition pre-close trades into \$500 dollar-volume buckets and compute the average absolute imbalance over the most recent 20 buckets:

$$\text{Trailing VPIN}_t = \frac{1}{20} \sum_{\tau=t-20}^{t-1} \frac{|V_\tau^B - V_\tau^S|}{V_\tau}$$

The question is whether this trailing measure, computed from data available at bucket $t-1$, predicts whether the maker loses money on trades in bucket t. Because every contract settles at \$1 or \$0, we can compute the maker's realized P&L on every bucket and test this directly.

The pooled panel contains 32,151 interior-price bucket observations across 534 single-name contracts and 544,142 interior-price observations across 17,298 broad-based contracts. Table 9 reports maker outcomes by quintile of trailing VPIN.

[Insert Table 9 about here]

[18] This is not a concern in equity markets, where prices do not converge to a known terminal value. It is a general feature of binary contract markets and would affect any flow-imbalance metric applied to prediction market data without correction.

[19] In the unrestricted sample, the YES-settlement rate among single-name bucket observations rises from 36% in the lowest trailing-VPIN quintile to 95% in the highest. In the interior sample, it ranges from 24% to 30%, confirming that the restriction eliminates the settlement-outcome confound.

On single-name contracts, the pattern is consistent with the Section 5 decomposition operating in real time. In the lowest trailing-VPIN quintile (Q1, mean = 0.54), the maker earns an average of \$41 per bucket and loses money on 43.1% of buckets. Profitability is roughly stable through Q4 (mean VPIN = 0.80, P&L = +\$46, loss rate 44.8%). In Q5 (mean VPIN = 0.90), maker P&L turns negative (−\$12) and the loss rate jumps to 51.7%. The YES-settlement rate varies only from 24% to 30% across quintiles, confirming that the interior restriction has eliminated the settlement-outcome confound.

On broad-based contracts, trailing VPIN has no predictive power at interior prices. Maker P&L is positive and roughly flat across all five quintiles (ranging from \$4 to \$11), and the loss rate barely moves (49.5% to 49.7%). The absence of predictive power is at first puzzling, given both the Panel A results for single-name contracts and VPIN's established track record in equity markets. But the resolution is clear once we consider the nature of the underlying events.

For VPIN to predict maker losses, the *direction* of one-sided flow must predict the *settlement outcome*, because it is the settlement outcome that determines whether the maker wins or loses. On single-name contracts, this condition holds. Among high-VPIN buckets on single-name contracts, those with YES-directed flow settle YES 36.4% of the time, while those with NO-directed flow settle YES only 25.8% of the time (logit $z = 12.30$, $p < 0.001$). The direction of one-sided flow carries information about whether the event will occur, and that information translates directly into maker losses when the flow is on the correct side.

On broad-based contracts, the condition fails. Among high-VPIN buckets, the YES-settlement rate is 50.4% when flow is YES-directed and 50.9% when flow is NO-directed, statistically indistinguishable (logit $z = -1.74$, $p = 0.08$). One-sided flow tells the maker nothing about which way the contract will settle. The reason is that broad-based contracts reference continuously moving underlyings. Consider a Bitcoin daily range contract trading at 50 cents. If Bitcoin rises in the preceding hour, takers buy YES, generating one-sided flow and high VPIN. But Bitcoin can just as easily reverse before settlement. The one-sided flow

reflects the *recent trajectory* of a continuously moving price, not private information about the *terminal value*. On a single-name contract such as "Will Netflix say 'Warner Bros'?", by contrast, one-sided NO flow during the earnings call reflects real-time observation that the phrase has not been uttered—information that is directly and irreversibly about the settlement outcome.[20]

6.3 Regression Analysis

To formalize these patterns, we estimate pooled bucket-level regressions of maker P&L on trailing VPIN, with standard errors clustered by contract:

$$\pi_{j,t}^{\text{maker}} = \alpha + \beta \cdot \text{Trailing VPIN}_{j,t} + \varepsilon_{j,t}$$

Table 10 reports the results for both OLS (dependent variable: dollar maker P&L on bucket t) and logit (dependent variable: indicator for maker P&L < 0).

[Insert Table 10 about here]

On single-name contracts, both specifications confirm the quintile pattern. Higher trailing VPIN predicts lower dollar P&L ($\hat{\beta} = -139$, $t = -2.42$, $p = 0.016$) and a higher probability that the maker loses on the next bucket (logit $\hat{\beta} = 0.93$, $z = 10.56$, $p < 0.001$). On broad-based contracts, neither specification is significant (OLS $t = 0.56$; logit $z = 0.65$), consistent with the absence of concentrated private information in macro and index contracts.

Because Section 5 established that the maker's profit is concentrated on NO-settling contracts, we can sharpen the test by estimating the regression separately within each settlement outcome. On NO-settling single-name contracts, where the maker's optimism-tax

[20] We confirm this interpretation by examining the distribution of maker P&L in high-VPIN buckets. On single-name contracts, the fraction of buckets with large maker losses (P&L $< -\$100$) rises from 35.3% in Q1 to 46.7% in Q5, while large wins (P&L $> +\$100$) remain roughly constant (46.3% to 45.8%). High VPIN on single-name contracts shifts the maker's outcome distribution toward losses. On broad-based contracts, both large wins and large losses increase proportionally in Q5 (47.9% and 47.2%, up from 39.5% and 38.5% in Q1). High VPIN on broad-based contracts increases the *variance* of maker outcomes without shifting the *mean*, because one-sided flow is equally likely to be on the winning or losing side of settlement.

profits are at risk from informed NO flow, trailing VPIN significantly predicts maker losses: OLS $\hat{\beta} = -160$ ($t = -1.99$, $p = 0.047$). On YES-settling single-name contracts, the coefficient is economically negligible and statistically insignificant ($t = -0.09$). The pattern is consistent with the mechanism identified in Section 5: informed traders erode maker profits specifically on NO-settling contracts by buying NO correctly, and trailing VPIN detects this erosion in real time.[21]

Sorting single-name bucket observations into finer bins reinforces the finding that the toxicity signal is concentrated at the extreme. Makers earn positive P&L at every level of trailing VPIN below approximately 0.86, but at the highest levels (VPIN above 0.94) maker P&L turns sharply negative (−\$43 per bucket, loss rate 55.3%). The maker's real-time risk from informed flow materializes not as a gradual erosion across the VPIN distribution but as a discrete jump at the upper tail.

These results complement the Kyle's λ findings from Section 4 through an independent, trade-based channel. Kyle's λ estimates the average price impact of order flow over a contract's lifetime. Trailing VPIN operates within a contract's life, identifying *when* the maker is absorbing toxic flow and predicting *which* specific buckets of trades are likely to generate losses. The two measures capture different facets of the same underlying phenomenon: informed order flow that the maker cannot distinguish from noise at the time of execution.

7 Discussion

7.1 The Behavioral Surplus as Adverse Selection Insurance

Wider spreads on single-name contracts confirm that makers recognize the adverse selection risk, but the binary payoff structure ensures that wider spreads cannot compensate for it:

[21] The asymmetry is not an artifact of the interior-price restriction. On YES-settling contracts, the maker is already expected to lose (the contract settles at 100 and the maker sold YES at a discount), so one-sided flow adds little incremental information about the maker's outcome. On NO-settling contracts, the maker expects to profit from uninformed YES takers, and one-sided NO flow signals that informed traders are eroding that expected profit.

the informed trader's edge can be orders of magnitude larger than the spread itself. Yet makers earn *more*—not less—on single-name contracts, because the behavioral surplus from YES-biased takers is concentrated precisely where the adverse selection is most severe. Single-name takers buy YES 60.3% of the time on contracts that settle YES only 30.9% of the time, generating a frequency edge (3.85 of 3.94 cents per trade) that more than absorbs the losses informed traders impose. The maker's entire profit comes from NO-settling contracts, where the optimism tax from uninformed YES takers exceeds the cost of trading against informed NO bettors by 1.52×. Trailing VPIN confirms that this erosion is detectable in real time: at interior prices, one-sided flow on single-name contracts predicts maker losses on the next volume bucket. This behavioral surplus has no close analog in equity markets, where the marginal trader is typically a sophisticated institution and the Glosten and Milgrom [1985] zero-profit condition binds through the spread.

7.2 Implications for Market Design and Regulation

Exchange operators face a fundamental tradeoff in choosing which contracts to list. Single-name contracts attract both informed traders (who impose adverse selection costs) and behaviorally motivated takers (who generate profits); broad-based contracts attract fewer of each. The net effect (maker profitability is *higher* on single-name contracts) suggests that the current equilibrium is stable. Makers have no reason to exit the very markets where adverse selection is most severe, because the behavioral surplus more than compensates.

This equilibrium, however, is fragile in a specific sense: it depends on the continued presence of uninformed, YES-biased takers. If the taker population became more sophisticated (through algorithmic trading tools, educational resources, or the exit of retail participants), the behavioral surplus would shrink and makers might face losses on high-adverse-selection contracts. Exchange operators would do well to monitor the taker informativeness rate (the fraction of takers who take the side that ultimately prevails) as an early-warning indicator of equilibrium erosion.

The Bezos contract demonstrates that informed trading occurs on prediction markets, and our cross-sectional results show that the contracts most susceptible to it (single-name contracts with concentrated private information) are identifiable ex ante. The combination of the taker-side field with settlement outcomes provides a surveillance tool of unusual power: regulators can compute taker correctness rates, identify contracts with abnormally informed flow, and flag accounts whose trading patterns are inconsistent with behavioral demand.

At the same time, the market-integrity case for aggressive enforcement is more nuanced here than in equities. Makers are not being driven from the market by informed trading. The behavioral surplus subsidizes their losses, maintaining liquidity even in the presence of insiders. This does not mean informed trading should be tolerated (it remains unfair to the uninformed takers who bear the losses), but it does mean that the standard market-quality justification for insider trading enforcement, that it deters market makers and reduces liquidity, may not apply with the same force in this setting. The prediction market maker's problem is not the Glosten and Milgrom [1985] zero-profit condition; it is whether the behavioral surplus remains large enough to absorb the information costs. That is a different problem, and it may call for different regulatory tools.

8 Conclusion

Adverse selection is the central challenge for any market that aggregates information through trading. Our paper asks how this challenge plays out in prediction markets, and what happens to the people who provide the liquidity that makes aggregation possible.

Using 41.6 million trades across 478,167 Kalshi contracts, we show that the standard tools of equity market microstructure (Kyle's λ, the Glosten-Harris decomposition, effective spreads) detect adverse selection in prediction markets precisely where theory predicts it. Kyle's λ in log-odds space is significantly elevated on single-name contracts ($t = 10.88$), the Glosten-Harris decomposition confirms that this price impact is almost entirely perma-

nent, and effective spreads are wider. But the binary payoff structure of prediction markets severs the link between wider spreads and maker compensation that is central to equity microstructure: the informed trader's potential edge (up to the full distance between the trade price and settlement) can exceed the spread by an order of magnitude, so the conventional spread-compensation mechanism breaks down.

The resolution lies in the maker's profit function. Single-name makers earn 3.94 cents per trade (over five times the 0.74 cents earned on broad-based contracts) despite facing more adverse selection. Nearly all of this profit (97.7%) comes from a frequency edge: takers buy YES 60.3% of the time on contracts that settle YES only 30.9% of the time. The behavioral surplus that uninformed takers generate is large enough to absorb the losses that informed traders impose, and it is concentrated precisely where the adverse selection is most severe.

A decomposition by settlement outcome sharpens the picture. The maker's entire profit comes from contracts that settle NO, where uninformed YES takers generate an optimism tax that exceeds the cost of trading against informed NO bettors by 1.52× on single-name and 1.28× on broad-based. Adapting VPIN to prediction markets as a real-time toxicity metric, we show that trailing VPIN at interior prices (30–70 cents, where a settlement-outcome endogeneity specific to binary contracts is absent) predicts maker losses on the next volume bucket for single-name contracts but not for broad-based. The effect is concentrated on NO-settling contracts, consistent with informed traders eroding the maker's optimism-tax profits in real time.

The broader implication is that adverse selection and market viability can coexist—indeed, can reinforce each other—through behavioral cross-subsidization. The maker's profit is a residual: the optimism tax from uninformed YES takers on NO-settling contracts minus the losses that informed NO traders impose. As long as the ratio exceeds one, the maker survives. This mechanism has no close analog in equity markets, where the marginal trader is typically a sophisticated institution. It may explain why prediction market liquidity has proven more resilient than standard microstructure theory would predict, and it suggests

that the health of these markets depends less on the severity of informed trading than on the continued willingness of uninformed takers to participate.

References

Arrow, K. J. (2008). The promise of prediction markets. *Science*, 320(5878), 877–878.

Becker, J. (2025). Microstructure of wealth transfer in prediction markets. Working paper, `jbecker.dev`.

Bürgi, C., Deng, C., and Whelan, K. (2026). Makers and takers: Economics of Kalshi. CEPR/CESifo Working Paper.

Copeland, T. E. and Galai, D. (1983). Information effects on the bid-ask spread. *Journal of Finance*, 38(5), 1457–1469.

Easley, D., López de Prado, M. M., and O'Hara, M. (2011). The microstructure of the "Flash Crash": Flow toxicity, liquidity crashes, and the probability of informed trading. *Journal of Portfolio Management*, 37(2), 118–128.

Easley, D., López de Prado, M. M., and O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies*, 25(5), 1457–1493.

Easley, D., López de Prado, M. M., and O'Hara, M. (2016). Discerning information from trade data. *Journal of Financial Economics*, 120(2), 269–286.

Easley, D. and O'Hara, M. (1987). Price, trade size, and information in securities markets. *Journal of Financial Economics*, 19(1), 69–90.

Easley, D., Kiefer, N. M., O'Hara, M., and Paperman, J. B. (1996). Liquidity, information, and infrequently traded stocks. *Journal of Finance*, 51(4), 1405–1436.

Butcher, S. (2025). Electronic market makers on Kalshi are small fish in a big pond. *eFinancialCareers*, `https://www.efinancialcareers.com/news/electronic-market-makers-trading-on-kalshi-are-small-fish-in-a-big-pond`.

Ellis, K., Michaely, R., and O'Hara, M. (2000). The accuracy of trade classification rules: Evidence from Nasdaq. *Journal of Financial and Quantitative Analysis*, 35(4), 529–551.

Glosten, L. R. and Harris, L. E. (1988). Estimating the components of the bid/ask spread. *Journal of Financial Economics*, 21(1), 123–142.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100.

Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, 46(1), 179–207.

Hasbrouck, J. (2004). *Empirical Market Microstructure*. Oxford University Press.

Hauser, F. and Rittler, D. (2024). Price formation in field prediction markets. *Journal of Financial Markets*, forthcoming.

Huang, R. D. and Stoll, H. R. (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3), 313–357.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.

Lee, C. M. C. and Ready, M. J. (1991). Inferring trade direction from intraday data. *Journal of Finance*, 46(2), 733–746.

Madhavan, A. (2002). Market microstructure: A practitioner's guide. *Financial Analysts Journal*, 58(5), 28–42.

Ng, J., Peng, L., Tao, R., and Zhou, G. (2026). Price discovery and trading in modern prediction markets. SSRN Working Paper No. 5331995.

O'Hara, M. (1995). *Market Microstructure Theory*. Blackwell, Cambridge, MA.

Palumbo, N. (2026). A microstructure perspective on prediction markets. SSRN Working Paper No. 6325658.

Swanson, E. T., Wang, C., and Wu, Y. (2025). Kalshi and the rise of macro markets. Federal Reserve Board Working Paper.

Tsang, K. P. and Yang, J. (2026). Political shocks and price discovery in prediction markets. arXiv:2603.03152.

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

Wolfers, J. and Zitzewitz, E. (2006). Interpreting prediction market prices as probabilities. NBER Working Paper No. 12200.

Figures and Tables

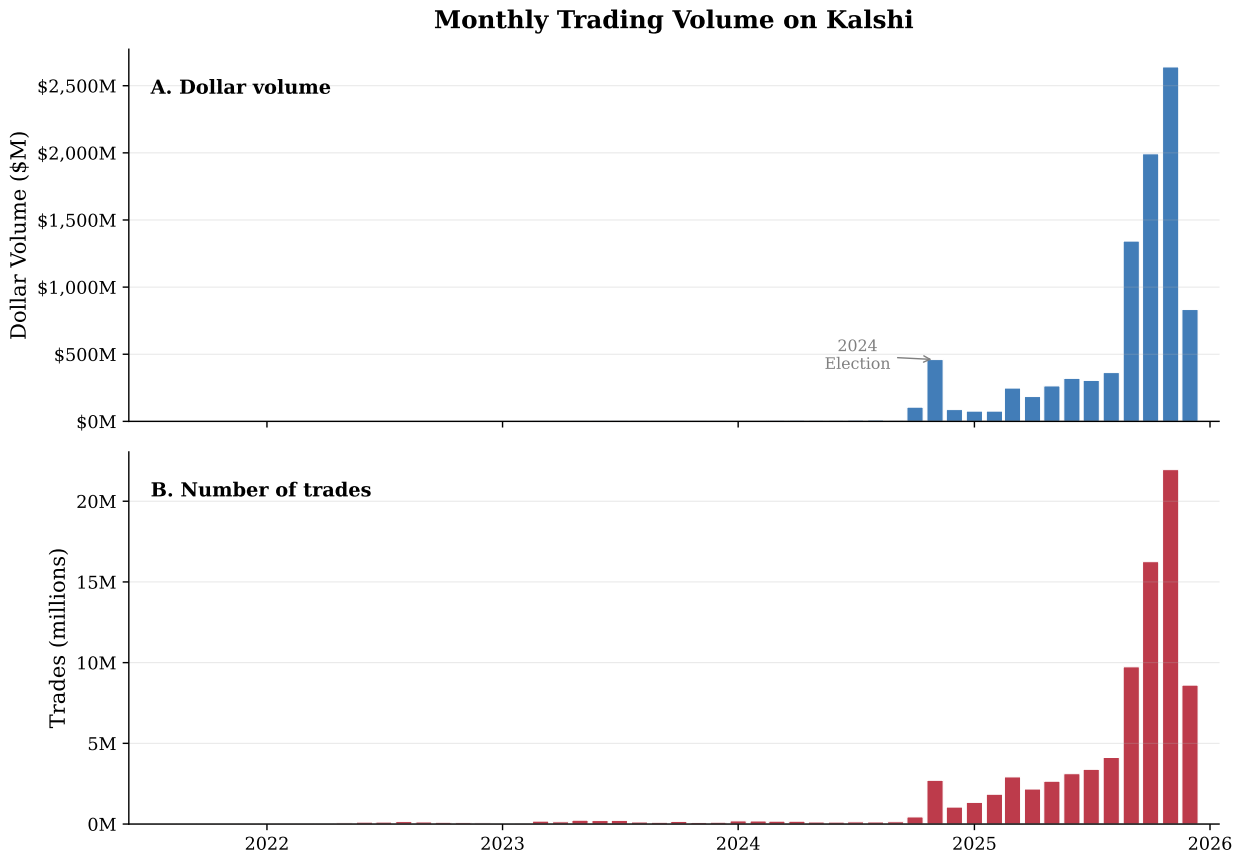


Figure 1: Monthly trading volume on Kalshi. Panel (a) shows dollar volume; Panel (b) shows number of trades. The exchange launched in mid-2021 and grew slowly through 2023, with volume accelerating sharply around the 2024 U.S. presidential election and continuing through the end of the sample in December 2025. Peak monthly dollar volume exceeded $2.6 billion in November 2025.

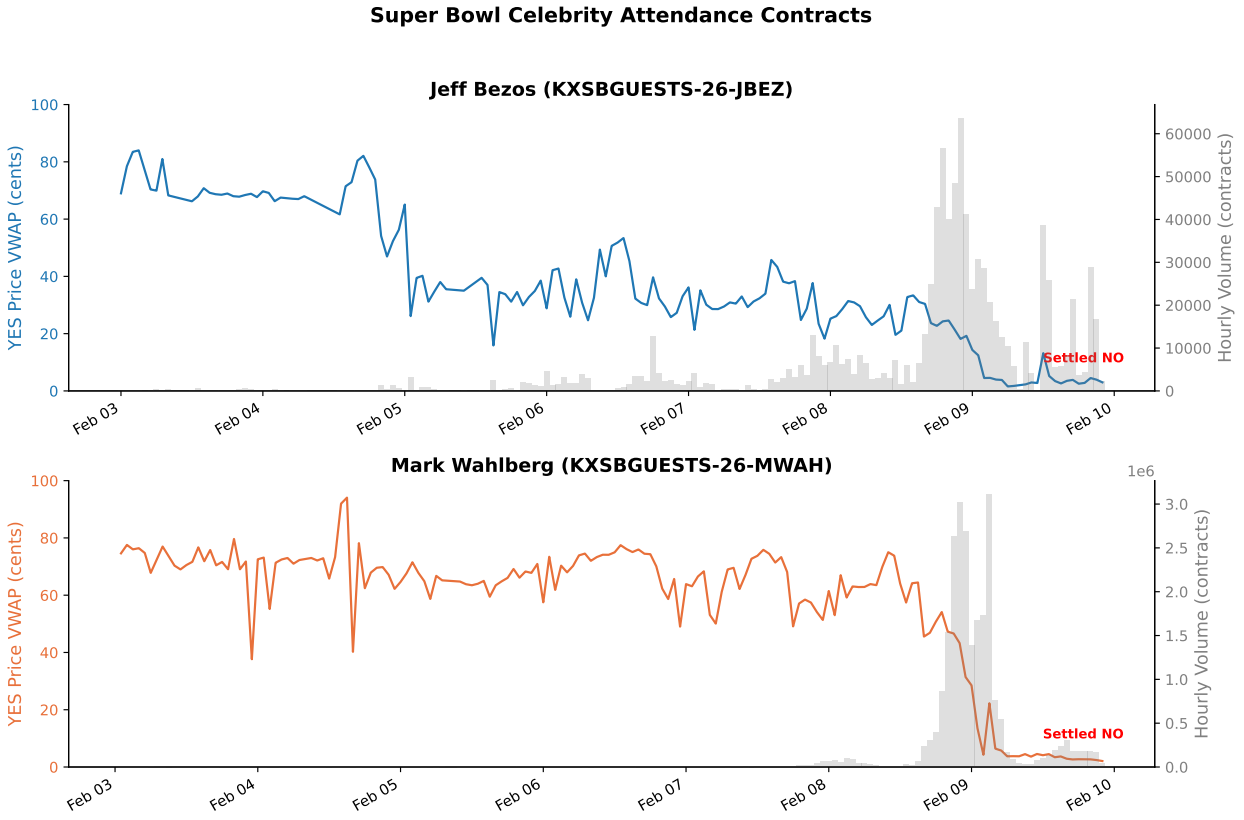


Figure 2: Price paths for two Super Bowl LX celebrity attendance contracts. *Top:* Jeff Bezos (KXSBGUESTS-26-JBEZ), where fraternity brothers with insider knowledge that Bezos would not attend drove the YES price from ∼70 cents down to ∼30 cents by buying NO. *Bottom:* Mark Wahlberg (KXSBGUESTS-26-MWAH), where a false rumor generated $24 million in volume. Both contracts settled NO. Gray bars show hourly volume.

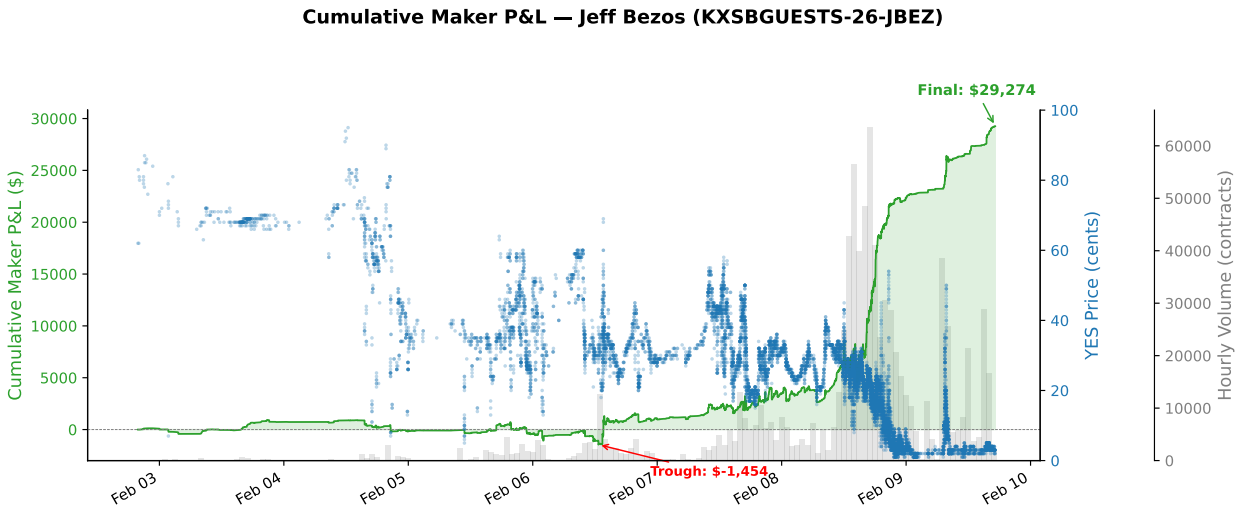


Figure 3: Cumulative maker P&L (green, left axis) and YES price (blue dashed, right axis) for the Bezos Super Bowl contract. Informed NO traders (fraternity brothers who knew Bezos would not attend) imposed periodic losses on makers throughout the contract's life. Yet the cumulative P&L remains positive because uninformed YES-biased takers (518K contracts) more than offset the informed NO flow (370K contracts). Final maker P&L: \$29,274.

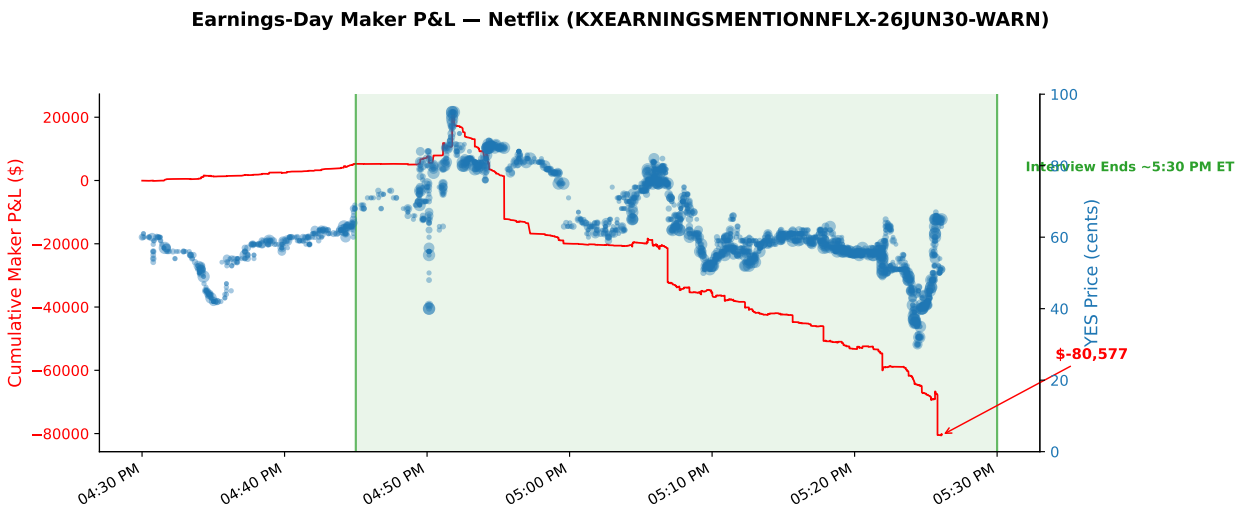


Figure 4: Cumulative maker P&L (red, left axis) and YES price (blue dots, right axis) for the Netflix "Warner Bros" earnings-mention contract. Panel (a) shows the full contract lifetime; Panel (b) zooms in on earnings day (January 20, 2026, 12:00–6:00 PM ET). Vertical lines mark the earnings release (4:00 PM), interview start (4:45 PM), and interview end (∼5:30 PM). During the live interview, NO-side takers (likely watching the broadcast) imposed $85,407 in losses on makers. This contract is the worst single-name outcome for makers in our sample.

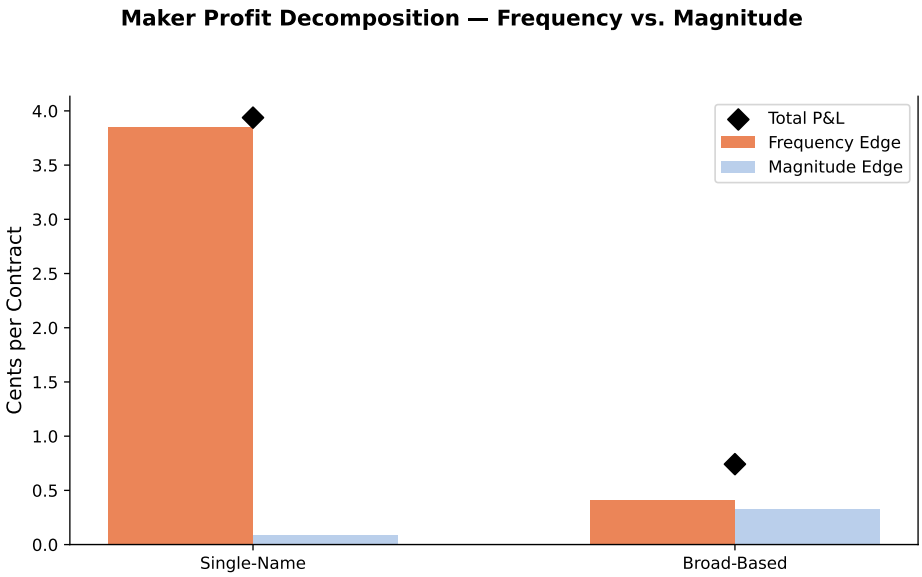


Figure 5: Decomposition of maker profits into frequency and magnitude components. Both contract types profit primarily through the frequency edge (winning more than 50% of trades), but the single-name frequency edge is an order of magnitude larger. Magnitude edges are near zero for both types. Diamonds mark total P&L per contract.

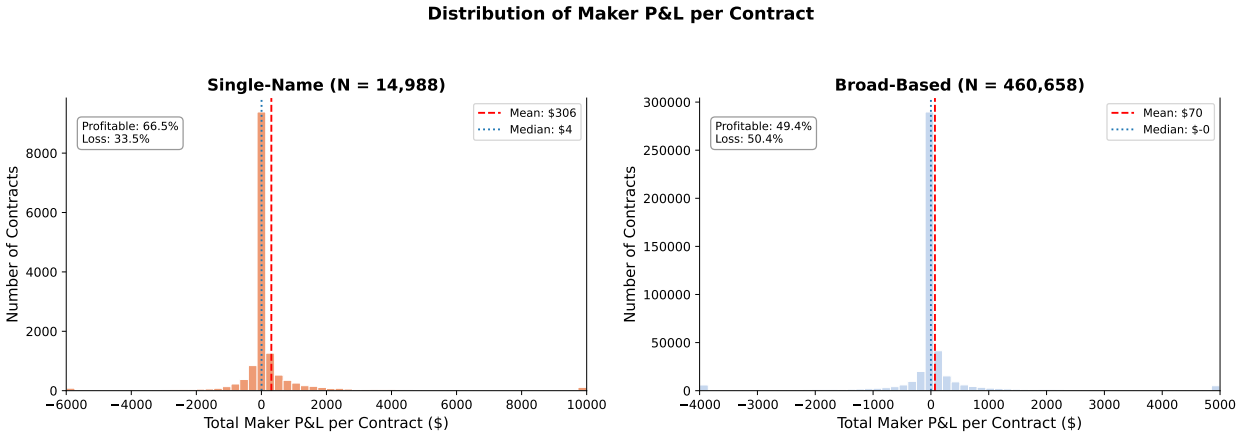


Figure 6: Distribution of total maker P&L across individual contracts. Single-name contracts (left, orange) and broad-based contracts (right, blue). Dashed vertical lines mark the median. Approximately 40% of contracts in both categories generate losses for the maker side, but right-skew in the profit distribution ensures positive average returns.

Table 1: Summary Statistics

	Single-Name	Broad-Based	Total
Markets	15,516	462,651	478,167
Settled markets	14,988	460,658	475,646
Total trades	3,022,734	38,572,352	41,595,086
Trade characteristics			
Mean yes_price	42.0¢	45.9¢	
Mean taker price (cost basis)	47.3¢	50.0¢	
Taker buys YES (%)	60.2%	53.2%	
Settlement			
Settles YES (%)	30.9%	38.4%	
Settles NO (%)	69.1%	61.6%	

Notes: The analysis sample consists of single-name and broad-based contracts with at least one recorded trade. Single-name contracts reference outcomes specific to a single company or individual; broad-based contracts reference macroeconomic indicators, market indices, or asset prices (see Appendix A for classification details). "Taker price" is the price paid by the taker: yes_price if taker buys YES, 100− yes_price if taker buys NO. Settlement rates are computed over settled markets with trades.

Table 2: Cross-Sectional Regressions: Adverse Selection on Contract Type

	(1) Kyle λ (raw)	(2) Kyle λ (log-odds)	(3) GH perm. (raw)	(4) GH perm. (log-odds)	(5) Eff. spread
SINGLE_NAME	0.0010***	0.0001***	0.0009***	0.0001***	0.2878***
	(9.20)	(10.88)	(7.46)	(9.58)	(3.17)
Controls	Yes	Yes	Yes	Yes	Yes
N	14,801	14,801	14,504	14,504	19,618
R^2	0.047	0.040	0.039	0.036	0.006

Notes: The unit of observation is the contract. Each column reports OLS estimates with HC1 standard errors. The dependent variable in columns (1)–(2) is Kyle's λ estimated per contract from 1-hour interval regressions; columns (3)–(4) use the permanent component from a Glosten-Harris decomposition; column (5) uses the contract-level mean of trade-level effective half-spreads. SINGLE_NAME equals 1 for contracts referencing outcomes specific to a single company or individual; the omitted category is broad-based contracts. Controls include log volume, price level, price level squared, and implied volatility. t-statistics in parentheses. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

Table 3: Robustness: Kyle's λ (Log-Odds) Under Alternative Specifications

	(1) No controls	(2) Financial only	(3) Insider conc.
SINGLE_NAME	0.0001***	0.0001***	
	(14.39)	(9.75)	
INSIDER_CONC			0.0002***
			(11.92)
Controls	No	Yes	Yes
N	14,801	13,842	14,801
R^2	0.020	0.035	0.044

Notes: Column (1) regresses log-odds λ on the single-name indicator alone. Column (2) restricts to financial contracts (company, index, macro, FX, commodity, crypto). Column (3) replaces the binary indicator with a continuous measure of insider information concentration (1.0 for single-name, 0.1–0.3 for broad-based based on the number of underlying entities). t-statistics in parentheses.

Table 4: Effective Spread and Midpoint Freshness

	OLS (equal-weighted)		WLS (volume-weighted)	
Window (τ)	β (SINGLE_NAME)	t	β (SINGLE_NAME)	t
$\leq$ 5 min	+0.289	1.60	+0.368	1.10
$\leq$ 10 min	+0.251	1.74	+0.445	1.46
$\leq$ 15 min	+0.217	1.79	+0.332	1.20
All ($\leq$ 60 min)	+0.173	2.27	+0.300	1.32

Notes: Each row restricts to trades occurring within τ minutes of the end-of-period candlestick quote and aggregates to the contract level. Contracts with fewer than 5 trades in the window are excluded. The dependent variable is the contract-level mean effective half-spread (cents). OLS gives each contract equal weight; WLS weights by number of trades. Controls include log volume, price level, price level squared, and implied volatility.

Table 5: Maker Profitability by Contract Type

	Single-Name	Broad-Based
Contracts (settled)	14,988	460,658
Trades	2,785,039	38,119,402
Per-trade		
Mean maker P&L/contract (cents)	3.94	0.74
Median maker P&L/contract (cents)	2.00	1.00
Win/loss		
Maker win rate (%)	56.6%	50.7%
Average win (cents)	29.24	31.80
Average loss (cents)	-29.06	-31.14
Per-contract		
% contracts maker profitable	66.5%	49.4%
Mean maker P&L/contract ($)	$305.89	$69.89
Median maker P&L/contract ($)	$4.48	$-0.01
Total maker P&L ($)	$4,584,642	$32,196,483

Notes: Maker P&L computed using realized settlement values. "Win rate" is the fraction of trades where the maker's side prevails at settlement. "Average win" and "Average loss" are the mean maker P&L per contract conditional on winning or losing. All amounts in cents per contract unless noted.

Table 6: Decomposition of Maker Profits

	Single-Name	Broad-Based
Total expected P&L	+3.94¢	+0.74¢
Frequency edge (Win rate − 50%) × (Avg win + Avg loss)	+3.85¢ (+6.6pp × 58.30¢)	+0.41¢ (+0.7pp × 62.94¢)
Magnitude edge 50% × (Avg win − Avg loss)	+0.09¢ (50% × 0.18¢)	+0.33¢ (50% × 0.66¢)

Notes: The frequency edge captures profits from winning more (or less) than 50% of trades; the magnitude edge captures profits from winning more per win than losing per loss. Amounts in cents per contract.

Table 7: Maker P&L by Taker Price Bucket (cents per contract)

	Single-Name		Broad-Based	
Taker price	Maker P&L	Maker wins	Maker P&L	Maker wins
1-10¢	+2.40¢	98%	+1.47¢	96%
11-20¢	+4.88¢	89%	+2.36¢	87%
21-30¢	+7.76¢	82%	+2.04¢	76%
31-40¢	+9.08¢	73%	+1.45¢	66%
41-50¢	+8.86¢	63%	+1.03¢	55%
51-60¢	+5.31¢	50%	+0.76¢	45%
61-70¢	+1.52¢	36%	−0.00¢	34%
71-80¢	−0.51¢	24%	−0.64¢	24%
81-90¢	+0.86¢	15%	−1.00¢	13%
91-100¢	+0.15¢	4%	−0.17¢	4%

Notes: Taker price is the cost basis: yes_price if taker buys YES, 100−yes_price if taker buys NO. Maker P&L is the average profit per contract (in cents) for trades where the taker paid within the given range. Maker wins is the fraction of trades where the maker's side prevails.

Table 8: Maker P&L by Settlement Outcome and Taker Side

		Single-Name		Broad-Based	
Settlement	Taker side	P&L ($M)	% of net	P&L ($M)	% of net
Panel A: By settlement outcome					
Settles NO		+6.25	+136%	+64.46	+200%
% contracts profitable		78.0%		54.8%	
Settles YES		−1.66	−36%	−32.27	−100%
% contracts profitable		40.8%		40.9%	
Net		**+4.58**		**+32.20**	
Panel B: By settlement outcome × taker side					
Settles NO	Taker buys YES	+18.22	+397%	+295.52	+918%
Settles NO	Taker buys NO	−11.97	−261%	−231.06	−718%
Settles YES	Taker buys YES	−8.31	−181%	−251.73	−782%
Settles YES	Taker buys NO	+6.65	+145%	+219.46	+682%
Panel C: Optimism tax vs. informed-trading cost (NO-settling)					
Optimism tax	(YES takers, settles NO)	+18.22		+295.52	
Informed cost	(NO takers, settles NO)	−11.97		−231.06	
Ratio		1.52×		1.28×	

Notes: All P&L figures are in millions of dollars. Panel A reports total maker P&L on settled contracts by settlement outcome. Panel B decomposes by both settlement outcome and taker side. "Taker buys YES" on a contract that settles NO is the optimism tax: the maker keeps the YES price because the event does not occur. "Taker buys NO" on a contract that settles NO is the informed-trading cost: the NO taker correctly bet against the event. Panel C reports the ratio of optimism-tax revenue to informed-trading cost on NO-settling contracts. Percentages express each cell as a fraction of the contract type's net maker P&L.

Table 9: Trailing VPIN and Next-Bucket Maker Outcomes (Interior Prices)

	Mean VPIN	Mean P&L	Median P&L	% Lost	% YES	n
Panel A: Single-Name Contracts (32,151 buckets, 534 contracts)						
Q1 (low)	0.54	+$41	+$74	43.1%	23.7%	6,431
Q2	0.67	+$36	+$78	44.8%	27.4%	6,430
Q3	0.74	+$42	+$91	44.7%	29.8%	6,430
Q4	0.80	+$46	+$99	44.8%	30.3%	6,430
Q5 (high)	0.90	−$12	−$44	51.7%	29.8%	6,430
Panel B: Broad-Based Contracts (544,142 buckets, 17,298 contracts)						
Q1 (low)	0.49	+$4	+$4	49.5%	51.5%	108,829
Q2	0.61	+$7	+$7	49.3%	51.4%	108,828
Q3	0.70	+$9	+$7	49.5%	50.0%	108,828
Q4	0.80	+$11	+$12	49.3%	47.9%	108,828
Q5 (high)	0.94	+$10	+$14	49.7%	50.1%	108,829

Notes: Each observation is a $500 dollar-volume bucket with volume-weighted average yes_price between 30 and 70 cents. Trailing VPIN is the mean absolute imbalance over the preceding 20 buckets. Mean and median maker P&L are in dollars per bucket. "% Lost" is the fraction of buckets where the maker's realized P&L is negative. "% YES" is the fraction of bucket observations on contracts that settle YES. Quintiles are formed within each contract type.

Table 10: Trailing VPIN and Next-Bucket Maker Outcomes: Regression Results (Interior Prices)

	OLS: Maker P&L ($)			Logit: Pr(Maker Loses)		
Subsample	$\hat{\beta}$	t	p	$\hat{\beta}$	z	p
Panel A: By contract type						
Single-name	−139	−2.42	0.016	0.93	10.56	<0.001
Broad-based	14	0.56	0.577	0.01	0.65	0.514
Panel B: Single-name, by settlement outcome						
Settles YES ($n = 9,067$)	−8	−0.09	0.924	0.36	2.11	0.035
Settles NO ($n = 23,084$)	−160	−1.99	0.047	1.00	9.71	<0.001

Notes: Panel A reports pooled bucket-level regressions of next-bucket maker P&L on trailing VPIN (20-bucket trailing window), restricted to buckets with VWAP between 30 and 70 cents. OLS standard errors are clustered by contract; logit uses HC1 standard errors. Panel B estimates the single-name regression separately within each settlement outcome. The dependent variable is maker P&L on bucket t (OLS, in dollars) or an indicator for maker P&L < 0 (logit).

Appendix A. Contract Classification

We begin with the full universe of over 7.3 million markets listed on Kalshi through early March 2026, encompassing 258.7 million trades across politics, sports, weather, entertainment, macroeconomic indicators, company-specific events, and other categories. We classify each market as single-name, broad-based, or unclassified based on a single economic criterion: how concentrated is private information about the contract's outcome?

Single-name contracts reference outcomes specific to a single company, individual, or small organization. A small number of corporate insiders, employees, or close associates possess decisive private information: whether Tesla will exceed a delivery target (known to Tesla executives), whether a specific word will be uttered during an earnings call (known to the speaker), or whether a celebrity will attend an event (known to the celebrity and their circle). The information structure is analogous to that of an individual stock.

Broad-based contracts reference macroeconomic aggregates, market indices, or asset prices where information is dispersed across thousands of analysts, traders, and policymakers. No individual or small group possesses a decisive private signal. Examples include S&P 500 index levels, CPI releases, Fed funds rate decisions, and cryptocurrency prices. The information structure is analogous to that of an index ETF or a macro futures contract.

Classification is performed using pattern-matching on the series ticker prefix assigned by Kalshi to each contract family. Applying these rules to the full universe yields 15,516 single-name markets and 462,651 broad-based markets with at least one recorded trade, which together constitute the analysis sample. The remaining markets (primarily politics, sports, weather, and entertainment) do not map cleanly onto either end of the information-concentration spectrum and are excluded from the main analysis. Politics contracts, for example, include both broad-based outcomes that aggregate millions of voter decisions and individual-action outcomes ("Will John Fetterman vote for RFK Jr.?") where a single decision-maker's intention is private information. Sports contracts involve competitive settings where coaches and players have asymmetric information about injuries and strategy,

but no individual can unilaterally determine the result. Because these categories blend elements of both concentrated and dispersed information, including them would blur the very distinction our empirical strategy is designed to exploit.

A.1 Single-Name Contracts

Table A.1 summarizes the composition of the single-name category.

[Insert Table A.1 about here]

Earnings call mentions dominate, accounting for the largest share of volume across over 100 publicly traded companies. These contracts are particularly well-suited for studying adverse selection: the outcome is determined during a discrete, observable event; company insiders and analysts who preview the call have a clear informational advantage; and each contract references a single corporate entity. Other significant subcategories include Spotify daily rankings (where employees at Spotify or music labels with access to real-time streaming data hold an informational edge), Tesla delivery and production targets, Super Bowl celebrity attendance, and various corporate events such as CEO departures, IPO timing, and M&A outcomes.

A.2 Broad-Based Contracts

Table A.2 summarizes the composition of the broad-based category.

[Insert Table A.2 about here]

Three subcategories dominate: Fed funds rate decisions, cryptocurrency (Bitcoin, Ethereum, and other digital assets), and equity indices (S&P 500 and Nasdaq 100 range contracts). All three share the defining characteristic that information about the outcome is widely dispersed: the federal funds rate is set by a 12-member committee based on publicly observed data, cryptocurrency prices are determined by a global decentralized market, and index levels aggregate information across hundreds of constituent stocks.

Table A.1: Single-Name Contract Composition

Subcategory	Markets	Series	Volume
Trump speech content	1,306	54	47,890,581
Earnings call mentions	2,635	107	44,151,438
Super Bowl attendance	41	2	40,270,790
Press conference mentions	1,441	3	28,828,987
Spotify daily rankings	5,642	22	26,388,967
Netflix rankings	2,003	12	22,937,244
Other (CEO changes, TikTok, etc.)	148	57	14,427,730
SpaceX launches/IPO	146	19	9,461,472
IPO timing (various)	345	21	8,395,714
Inauguration attendance	36	3	6,636,894
App store rankings	1,434	2	4,855,665
Tesla deliveries/production	181	33	4,576,132
M&A outcomes	3	1	2,733,515
Musk-specific (tweets, lawsuit)	114	3	1,644,286
OpenAI-specific	41	12	1,019,166
Total	**15,516**	**351**	**264,218,581**

Notes: Volume is in contracts.

Table A.2: Broad-Based Contract Composition

Subcategory	Markets	Series	Volume
Cryptocurrency (BTC, ETH, etc.)	303,336	116	2,564,137,697
Fed funds / rate decisions	1,198	33	657,890,989
Equity indices (S&P 500, Nasdaq)	86,516	38	569,082,997
Recession / gas prices / debt	605	30	99,801,263
CPI / inflation	2,392	29	61,882,065
Employment / labor	1,091	11	25,293,404
Commodities / rates (oil, gold, Treasuries)	5,112	22	22,335,847
Foreign exchange (USD/JPY, EUR/USD)	61,940	17	19,899,880
GDP	295	26	16,498,151
PCE / other macro	166	4	852,608
Total	**462,651**	**326**	**4,037,674,901**

Notes: Volume is in contracts.

Appendix B. Informed Trading Across the Full Kalshi Universe

The main analysis restricts attention to single-name and broad-based contracts, where we hypothesize that the cross-sectional variation in information concentration is sharpest. This appendix extends the analysis to the full universe of Kalshi contracts. The goal is to assess whether the single-name/broad-based dichotomy captures the extremes of informed trading or whether other contract categories exhibit comparable or greater levels of adverse selection.

B.1 The Kalshi Contract Universe

Kalshi lists over 7.3 million markets spanning company-specific events, macroeconomic indicators, politics, sports, entertainment, weather, and other categories. Table B.1 summarizes the universe by contract type.

Table B.1: Kalshi Contract Universe

Category	Contracts	% of Total	Mean Trades
Single-name	15,516	0.2%	17,029
Broad-based	462,730	6.3%	8,727
Politics	9,856	0.1%	191,119
Sports	340,973	4.6%	97,349
Entertainment	11,614	0.2%	30,339
Weather	45,658	0.6%	12,434
Other	94,357	1.3%	71,190
Parlay (multi-event)	6,385,199	86.7%	633
Total	7,365,903	100%	6,894

Notes: Contract counts reflect all markets with at least one recorded trade through March 11, 2026. Mean trades is the average number of trades per contract within each category. Single-name and broad-based are defined as in Appendix A. Remaining categories are assigned using series-ticker pattern matching on the full Kalshi taxonomy.

Two features of this table warrant discussion. First, parlay contracts account for 86.7% of all listed markets. These are programmatically generated multi-leg contracts that bundle two

or more standalone propositions into a single bet. Consider a Buffalo Bills game against the Miami Dolphins. Kalshi offers standalone contracts on the game's outcome (Will the Bills win?), the total points (Over 47.5?), and individual player statistics (Will Josh Allen throw for over 275.5 yards?). Kalshi also generates every possible combination of these propositions as a separate contract: "Bills win AND over 47.5 points" is one ticker; "Bills win AND Josh Allen over 275.5 yards AND Tyreek Hill scores a touchdown" is another. With 10 standalone propositions on a single game, the number of possible multi-leg combinations runs into the hundreds; across a full season of games, the combinatorics produce millions of tickers. These contracts are not dormant (95% have volume of at least 10), but they are compositional of the standalone contracts that form the basis of our analysis. Because their information structure is inherited from the component legs rather than generated independently, we exclude them throughout.

Second, politics contracts are the most actively traded per contract (191,119 mean trades), driven by presidential election markets that attracted billions of dollars in volume during the 2024 cycle. Sports contracts are next, followed by other and entertainment. Single-name contracts, while few in number, are moderately active (17,029 mean trades), providing sufficient data for reliable estimation.

B.2 Kyle's λ Across All Categories

We replicate the Table 2 specification on the full sample of non-parlay contracts, with single-name as the omitted category:

$$\hat{\lambda}_j = \beta_0 + \sum_{k \neq \text{SN}} \beta_k \cdot \mathbf{1}[\text{Category}_j = k] + \gamma' X_j + \varepsilon_j$$

where X_j includes the same controls as in the main analysis (log volume, price level, price level squared, and implied volatility). We also report a specification with log market duration as an additional control, since contract duration varies substantially across categories

(weather contracts have a median duration of 39 hours; politics contracts, 927 hours). All results include heteroskedasticity-robust standard errors. The sample comprises 129,731 contracts with at least 50 five-minute trading intervals.

Table B.2 reports the results.

Table B.2: Kyle's λ (Log-Odds) by Contract Category

	(1) Baseline	(2) + log(duration)
BROAD_BASED	-0.000106^{***}	-0.000111^{***}
	(-11.28)	(-11.83)
POLITICS	-0.000073^{***}	-0.000058^{***}
	(-6.78)	(-5.31)
SPORTS	-0.000096^{***}	-0.000102^{***}
	(-10.97)	(-11.69)
ENTERTAINMENT	-0.000117^{***}	-0.000107^{***}
	(-11.15)	(-10.23)
WEATHER	0.000070^{***}	0.000057^{***}
	(7.36)	(5.89)
OTHER	-0.000061^{***}	-0.000065^{***}
	(-6.77)	(-7.20)
Controls	Yes	Yes
N	129,731	129,731
R^2	0.101	0.102

Notes: OLS with HC1 standard errors; t-statistics in parentheses. Omitted category: single-name (15,516 contracts classified as in Appendix A). Multi-event (parlay) contracts excluded. Controls: log volume, price level, price level^2, implied volatility. Column (2) adds log(market duration), where duration = close_time − open_time. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Every category exhibits significantly lower Kyle λ than single-name contracts, with one exception: weather. Broad-based contracts ($\beta = -0.000106$, $t = -11.28$), sports ($t = -10.97$), entertainment ($t = -11.15$), politics ($t = -6.78$), and other ($t = -6.77$) all have substantially lower price impact of order flow. The magnitudes are large: the broad-based coefficient represents a reduction of roughly half relative to the single-name mean. Weather contracts, however, exhibit significantly *higher* λ than single-name ($\beta = 0.000070$, $t = 7.36$), and this effect persists after controlling for market duration ($\beta = 0.000057$,

$t = 5.89$).

The weather result is puzzling. Who has a decisive private signal about tomorrow's high temperature in Chicago? The next section resolves this puzzle.

B.3 Directional Lambda: Separating Informed from Uninformed Price Impact

Kyle's λ measures the magnitude of price response to order flow, but it does not distinguish whether that order flow is on the correct side of the eventual settlement outcome. In equity markets, this distinction is unavailable because the terminal value is unobserved. In prediction markets, where every contract settles at 0 or 100, we can make it directly.

For each contract j with settlement value $V_j \in \{0, 100\}$, we classify each 5-minute interval h based on whether signed order flow is on the correct side of settlement:

$$\text{CorrectSide}_{j,h} = \mathbf{1}\big[\text{sign}(\text{SOF}_{j,h}) = \text{sign}(V_j - 50)\big]$$

Net buying (SOF > 0) is correct-side when $V_j = 100$ (the contract settles YES), and net selling (SOF < 0) is correct-side when $V_j = 0$. Intervals with zero net flow are excluded. We then estimate λ separately on each subset:

$$\text{Correct-side:} \quad \Delta\ell_{j,h} = \alpha_j^C + \lambda_j^C \cdot \text{SOF}_{j,h} + \varepsilon_{j,h}^C \quad \text{for } h \in \{\text{CorrectSide} = 1\} \tag{1}$$

$$\text{Wrong-side:} \quad \Delta\ell_{j,h} = \alpha_j^W + \lambda_j^W \cdot \text{SOF}_{j,h} + \varepsilon_{j,h}^W \quad \text{for } h \in \{\text{CorrectSide} = 0\} \tag{2}$$

We require at least 25 intervals per subset. The directional difference $\lambda_j^D = \lambda_j^C - \lambda_j^W$ captures whether order flow has more price impact when it is on the informationally correct side. A positive λ^D is the signature of genuine informed trading: prices move more when the flow knows something.

A design note: we split on the direction of the *regressor* (SOF relative to settlement)

rather than on the direction of the *dependent variable* (the price change). Splitting on the dependent variable would introduce selection bias, since intervals with large $|\Delta\ell|$ would be mechanically more likely to be classified as toward-settlement.

Table B.3 reports cross-sectional regressions of each component on category dummies, with single-name as the omitted group.

Table B.3: Directional Lambda: Correct-Side vs. Wrong-Side Price Impact

	(1) λ^C (correct-side)	(2) λ^W (wrong-side)	(3) $\lambda^C - \lambda^W$
BROAD_BASED	-0.000382^{***}	-0.000004	-0.000354^{***}
	(-20.90)	(-0.27)	(-16.11)
POLITICS	-0.000258^{***}	0.000040^{***}	-0.000237^{***}
	(-12.03)	(2.61)	(-8.91)
SPORTS	-0.000299^{***}	-0.000020^{*}	-0.000211^{***}
	(-16.15)	(-1.66)	(-9.34)
ENTERTAINMENT	-0.000366^{***}	0.000001	-0.000309^{***}
	(-18.33)	(0.07)	(-12.74)
WEATHER	-0.000044^{**}	0.000103^{***}	-0.000121^{***}
	(-2.22)	(7.50)	(-4.94)
OTHER	-0.000195^{***}	-0.000008	-0.000072^{***}
	(-10.13)	(-0.60)	(-2.98)
Controls	Yes	Yes	Yes
N	85,832	98,459	60,133
R^2	0.098	0.015	0.053

Notes: OLS with HC1 standard errors; t-statistics in parentheses. Omitted category: single-name. Controls: log volume, price level, price level^2, implied volatility, log(market duration). Sample restricted to settled non-parlay contracts with ≥ 50 total intervals and ≥ 25 intervals per directional subset. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The decomposition resolves the weather puzzle. Column (1) reports the price impact of order flow on the correct side of settlement. On this dimension, single-name contracts dominate every category: broad-based ($t = -20.90$), entertainment ($t = -18.33$), sports ($t = -16.15$), politics ($t = -12.03$), other ($t = -10.13$), and weather ($t = -2.22$) all exhibit significantly lower correct-side λ. When order flow is on the right side of the settlement outcome, single-name contracts show the strongest price response.

Column (2) reveals the source of weather's anomalous overall λ: wrong-side price impact.

Weather contracts have the highest wrong-side λ of any category ($\beta = 0.000103$, $t = 7.50$), meaning that prices on weather contracts respond more strongly to order flow that turns out to be on the losing side of settlement than on any other type of contract. Prices move, but the flow is wrong. Politics also shows a positive wrong-side coefficient ($t = 2.61$), but the effect is smaller. For broad-based, sports, entertainment, and other contracts, wrong-side price impact is statistically indistinguishable from single-name.

Column (3) reports the directional difference $\lambda^C - \lambda^W$. Every category has a significantly negative coefficient: the gap between correct-side and wrong-side price impact is largest for single-name contracts. This is exactly what informed trading should produce. When insiders trade on single-name contracts, prices respond; when noise traders push prices in the wrong direction, the response is muted. On weather contracts, the pattern is reversed: wrong-side flow moves prices more than correct-side flow, consistent with uninformed or contrarian trading that happens to generate large price impact without carrying information about the terminal value.

The overall picture is consistent with the main analysis. The single-name/broad-based dichotomy captures the two extremes of informed trading on Kalshi. Single-name contracts exhibit the highest price impact from order flow, and that price impact is concentrated on the correct side of the settlement outcome. Broad-based contracts sit at the opposite end, with the lowest correct-side λ and a negligible directional gap. The remaining categories fall between these poles, with weather providing an instructive case in which high overall λ reflects uninformed price impact rather than genuine information.