Journal of Quantitative Description: Digital Media 6(2026), 1–92 10.51685/jqd.2026.011

Electoral Predictions on Polymarket:

A Quantitative Description of Trading, Commenting, and Reacting

Ross Dahlke, Naomi Mine, Haotian Zhao, Yinqi Huang, Dhavan Shah

University of Wisconsin–Madison, School of Journalism and Mass Communication, USA

Election betting platforms, which combine real-money trading with public political discussion and social feedback, have seen explosive growth worldwide. Prediction-market research emphasizes aggregate accuracy and price dynamics, while computational political communication research examines discourse in digital forums; however, these literatures remain analytically separate. Here, we provide a quantitative description of how financial, discursive, and social behaviors are jointly observable on a single platform. We analyze 778,634 accounts (unique wallet addresses) over 846 days, documenting 30,502,864 trades totaling $5.86 billion, 384,586 comments, and 571,523 reactions across three behavioral modalities (trading, commenting, reacting) on the Polymarket platform, spanning 959 election-related events. We find that within events, trading and commenting are largely decoupled and temporally ordered: among wallets exhibiting both behaviors, trading and commenting co-occur in only 21.8% of wallet-event combinations, are only

Author contact: ross.dahlke@wisc.edu

moderately correlated within events, and 90.2% of joint pairs are trade-first. Participation is highly concentrated across all behaviors, with concentration varying systematically by behavior type: financial trading (Gini = 0.951; top 1% account for 73.7% of volume) concentrates most sharply, followed by social reactions (Gini = 0.845), then political discourse (Gini = 0.814; top 10% produce 76.4% of comments). This cross-behavior ordering persists when the analysis is restricted to higher minimum-activity thresholds. We also document the volume of single-action wallets: 12.6% of trading accounts execute exactly one trade, 41.1% of commenters post exactly one comment, and 34.3% of reactors leave exactly one reaction. While single-mode accounts (engaging in only one modality) dominate numerically (95.9% of accounts, predominantly Trader Only at 94.4%), a small multi-modal minority (engaging in two or more modalities; 4.1%) contributes disproportionately across financial, discursive, and social behaviors. These findings characterize observable wallet-level political activity when financial commitment is visible alongside discursive acts, informing comparative and explanatory research on hybrid platforms. More broadly, quantifying who participates, how activity concentrates, and when engagement occurs provides a foundation for studying participation across behavioral domains that have typically been analyzed separately.

Keywords: *Polymarket, Political communication, Visibility affordances, Multimodal participation, Participation inequality*

Introduction

Election betting platforms combine real-money trading with public political discussion and visible social feedback, yet research rarely describes these behaviors together on the same platform. Dating back to the 19th century, when betting markets accurately predicted election outcomes before scientific polling (Rhode & Strumpf, 2004), these new platforms saw unprecedented adoption during the 2024 U.S. Presidential Election, with over $3.7 billion in bets on Polymarket alone for the 2024 U.S. Presidential Election market.[1]

These spaces create an emerging class of hybrid digital environments where observable political action may carry both discursive and financial consequences. Unlike traditional social media, where talk is reputational but low-cost, or private betting, where stakes remain hidden, modern prediction platforms make both arguments and financial positions publicly observable and verifiable. As these systems attract broader participation and regulatory attention, systematically describing who participates, how activity concentrates, and how financial and discursive behaviors relate helps characterize voice and influence in hybrid political-financial spaces.

In this study, we document who engages across trading, commenting, and reacting; how

[1] Polymarket Presidential Election Winner 2024 market: https://polymarket.com/event/presidential-election-winner-2024?tid=1762042985009, accessed November 2025. Trading on Polymarket was prohibited for U.S.-based users under CFTC regulations during this period, though VPN access remained technically possible. Platform data do not reveal user geography.

activity concentrates within each modality; how behaviors associate among accounts engaging in multiple modes; and when engagement varies relative to electoral events. Using publicly visible digital traces—trading histories and positions, comment threads, and reaction patterns—we document 778,634 accounts across 846 days (May 2023–September 2025), analyzing 30,502,864 trades totaling $5.86 billion, 384,586 comments, and 571,523 reactions from 959 election-related events. We conduct this descriptive research because systematic documentation provides the empirical foundation needed to understand emerging communication phenomena, particularly in rapidly evolving digital environments where assumptions quickly become outdated (Gerring, 2012; Karpf, 2024). Political betting platforms represent such contexts: regulatory frameworks shifted during the 2024 election cycle, user bases expanded from specialized traders to broader participation, and platform affordances integrated financial and discursive behaviors in ways not present in historical betting markets. Without systematic documentation of participation patterns, theoretical explanations risk detachment from observable account behaviors (Munger et al., 2021).

To address this descriptive gap, we document account-level participation across three behavioral modalities (trading, commenting, and reacting) organized into five research dimensions: account composition and cross-mode distribution (RQ1); financial engagement and trading concentration (RQ2); political discourse and comment production (RQ3); social feedback through reactions (RQ4); and within-event cross-modal participation (RQ5). We distinguish *single-mode* accounts (engaging in only one modality) from *multi-modal* accounts (engaging in two or more). We treat this distinction as a descriptive partition of accounts by presence in modalities rather than as a substantive claim about behavioral types.

We reveal three key patterns. First, single-mode accounts dominate numerically (95.9% of accounts) while a small multi-modal minority (4.1%) drives disproportionate activity, and within-event trading and commenting actually co-occur in only 21.8% of wallet-event combinations among wallets active in both modalities, with trading almost always first when they do (90.2%; consistent with the wallet-aggregate Pearson $r = 0.28$, 95% CI [0.27, 0.29]), indicating sparse and temporally ordered cross-modal engagement, suggesting distinct participation logics within accounts. Second, concentration is extreme but varies systematically, reflecting two distinct dimensions of participation: an extensive margin dominated by low-activity participants (most traders execute only a small number of trades, most commenters post once or twice, most reactors leave a single reaction) and a smaller intensive margin producing the bulk of activity. Across modalities, financial behaviors (Gini = 0.951; top 1% account for 73.7% of volume) concentrate most sharply, followed by reactions (Gini = 0.845), then discourse (Gini = 0.814; top 1% produce 36.0% of comments). We report each headline statistic for the full population and across sensitivity analyses that vary the minimum-activity threshold, so readers can see how the statistic moves as low-activity accounts are progressively excluded. Third, engagement varies temporally across the electoral cycle, with trading composition shifting toward single-mode traders during high-stakes periods while commenting exhibits event-driven spikes. For political communication and computer-mediated communication scholarship, these findings provide a descriptive baseline for research on hybrid financial–discursive environments in which visibility, participation inequality, and cross-mode dynamics can be studied together.

Background and Related Work

Financial Political Engagement and Information Aggregation

Political betting markets have a rich history of using financial stakes to aggregate political information. Markets with monetary incentives have consistently outperformed traditional polling mechanisms in electoral prediction accuracy: from 1884 to 1940, the mid-October betting favorite won 11 of 12 decided U.S. presidential elections, with up to $165 million (in 2002 dollars) wagered in a single race (Rhode & Strumpf, 2004), and pre-1936 market prices performed nearly as well as modern poll-based forecasts (Erikson & Wlezien, 2012). Contemporary prediction markets have demonstrated accuracy across electoral and non-electoral domains (Diercks et al., 2026; Johansson, 2024; Rabetti et al., 2026), though platform-era evidence shows accuracy varies across exchanges (Clinton & Huang, 2025). Real-time order-flow signals have also flagged geopolitical shocks in advance of official announcements: volatility indices summarizing Kalshi and Polymarket price activity detected the February 2026 Israel-Iran strike roughly 20 minutes before the Israeli Defense Forces announcement and approximately an hour before mainstream news coverage (Brown, 2026).

Political betting traditions persist across democratic systems with varying regulatory frameworks, demonstrating the cross-cultural robustness of financial political expression (Rhode & Strumpf, 2008). These historical precedents position political betting as a meaningful form of political engagement. Meta-analytic evidence documents positive associations between digital media use and political participation, with effects increasing over time as platforms mature and user bases diversify (Boulianne, 2015, 2020); betting platforms extend this pattern by combining financial participation with the discursive behaviors that meta-analyses have typically studied in isolation.

Modern digital betting platforms extend this participatory framework through what recent scholarship terms *political investorism*: "the individual or collective use of a financial stake to express political values" (O'Brien et al., 2023). Investorism is conceptually distinct from political consumerism: investors occupy an ownership rather than consumption position, pursue profit-seeking rather than need-meeting, and deploy capital repeatedly with positional rights rather than via one-time transactions (O'Brien et al., 2023), and survey-based factor analyses identify investorism as a separate latent dimension of political participation (Christensen & Brännlund, 2024). Betting market accounts extend this framework further by holding observable positions that function simultaneously as speculative instruments and political signals: financial positions may operate as observable political signals, though they do not reveal preferences one-to-one and may also reflect sophisticated arbitrage strategies (Chen et al., 2025; Jain et al., 2025), while repeated engagement also shapes attitudes, with experimental findings showing that stock trading shifts socioeconomic values and policy preferences (Margalit & Shayo, 2021). This bidirectionality reflects the integration of identity and economic decision-making (Akerlof & Kranton, 2000) and manifests in partisan-driven allocation that may defy rational expectations of returns. Unlike ethical investing or shareholder activism, betting market investorism entangles profit-seeking with political identity expression, with platform affordances making both financial and discursive acts visible and consequential.

Web3-based platforms like Polymarket now provide open digital trace data enabling large-scale quantitative description of these patterns for the first time, and given the necessity of empirical description before theorizing in rapidly evolving digital contexts (Gerring, 2012; Karpf, 2024), particularly where systematic account-level documentation has been absent (Munger et al., 2021), we ask, *(RQ1) How do accounts distribute across and within participation modalities in election betting markets, and how does engagement concentrate and vary over time?*

Trading Behavior and Platform Affordances

Platform affordances structure behavior through design choices that make specific actions observable and accessible, with visibility mechanisms shaping what participants perceive and do (Treem et al., 2020). Political betting platforms introduce a distinct affordance within political communication spaces: account-level visibility and access to financial trading in political outcomes. Financial positions (trades, holdings, and realized gains or losses) are publicly surfaced at the account level, and design features make individual financial activity visible to other participants, making trading available as public evidence that others may interpret as political conviction (Preda, 2013). These affordances render financial behavior socially situated. Furthermore, research on prediction-market trading documents heterogeneous outcomes across price ranges and trader roles (Bartlett & O'Hara, 2026; Bürgi et al., 2025; Deleep et al., 2026; Whelan, 2023, 2025), with complementary evidence reframing this variation as a function of trader-level skill and execution rather than role per se (Della Vedova, 2026; Yang, 2026). Yet, a systematic account-level description of how these visible financial affordances shape political behavior remains limited, despite trading functioning as the platform's central affordance, raising descriptive questions about who participates, how participation distributes across accounts, and how it aligns with high-salience political moments. To document these patterns, we ask: *(RQ2) How does financial trading activity distribute, concentrate, and vary temporally across account types in election betting markets?*

Communicative Affordances in Political-Financial Contexts

Political betting platforms serve as sites of electoral discussion and political communication, where participants engage in sustained discourse on electoral outcomes, candidate viability, polling data, and campaign developments. Political discussion occurs extensively in digital spaces not primarily designed for political discourse (Rajadesingan et al., 2021), often emerging as incidental byproducts of non-political activities (Minozzi et al., 2020). These spaces constitute venues for political dialogue distinct from traditional political forums: participants discuss electoral politics while holding financial stakes in the outcomes they debate, creating environments where political argumentation intersects with monetary commitment. Here, we provide an empirical description of who participates in political communication within betting contexts, how discourse production concentrates across accounts, what temporal patterns structure discursive engagement relative to electoral events, and how financial and discursive participation intersect among accounts active in both modes.

Besides functioning as political discussion venues, platform affordances structure political expression through design choices that make communication visible to others in ways previously impractical in historical, non-digital betting markets (Treem et al., 2020). Communication visibility creates multidimensional effects: making discourse observable to audiences, preserving communicative acts over time, enabling asynchronous participation, and generating persistent records subject to evaluation and reinterpretation. Mediation processes, rather than technological features alone, are the appropriate analytical focus for understanding communicative behavior (Carr, 2020), directing analytical attention to how platform processes shape individual communication patterns and strategic choices. In financial contexts specifically, internet stock message boards demonstrate that online discussion predicts trading volume and volatility beyond traditional media effects (Antweiler & Frank, 2004), while social-media-based prediction markets elicit voluntary discursive contributions alongside financial participation (Qiu & Kumar, 2017). These findings suggest that visibility mechanisms operate differently in contexts where discourse carries financial stakes, creating environments in which communication serves simultaneously informational, strategic, and social functions.

Given (i) political betting platforms as understudied sites of electoral discussion and political communication, (ii) these spaces enabling sustained political dialogue where financial stakes intersect with political argumentation, and (iii) the absence of systematic empirical description of account-level participation in these hybrid political communication environments, we examine how discourse unfolds alongside publicly visible financial positions. We ask: *(RQ3) How does commenting activity distribute, concentrate, and associate with trading behavior across account types, and how do temporal patterns of political discourse relate to electoral events?*

Beyond discursive production, political expression also operates through reactive signaling affordances. While comments require compositional effort and expose accounts to challenges of argumentation, reactions provide low-cost mechanisms for social feedback and affiliation. Reactions vary systematically by political ideology, with nationalist and populist parties eliciting higher proportions of angry reactions across national contexts (Muraoka et al., 2021), suggesting that micro-interactions reflect substantive political orientations. Content curation choices are driven by affiliative motivations to connect with politically like-minded audiences (Marie & Petersen, 2025), while small-scale sanctioning behaviors serve as identity-signaling and norm-enforcement mechanisms (Ahn et al., 2024). These reactive signals provide social feedback with minimal compositional cost, functioning as complementary affordances to discursive commentary within visible platforms.

However, systematic empirical description of how accounts employ reactive affordances in contexts combining political expression with financial stakes remains limited. Here, we provide account-level documentation of reaction-giving patterns, concentration of reactive engagement across accounts, and temporal patterns of social signaling within political-financial environments. Because reactions are low-cost, visible signals that can express affiliation or sanctioning alongside financial positions and political discourse, we ask: *(RQ4) How does social reaction activity distribute, concentrate, and vary temporally across account types in election betting markets?*

Given that online discussion is associated with trading volume and volatility in financial settings (Antweiler & Frank, 2004), and prediction-market environments elicit voluntary discursive contributions alongside financial participation (Qiu & Kumar, 2017), both of which raise the question of whether trading and commenting actually co-occur at the event level and how they are temporally ordered when they do. This event-level lens prevents aggregate co-presence from being mistaken for engagement around the same electoral object. To document these patterns, we ask, *(RQ5) When wallets active across multiple modalities engage with the same political event, do trading and commenting actually co-occur within the event, how are they temporally ordered when they do, and how does within-event coupling compare to wallet-aggregate measures?*

In the present study, we use the term *modality* as descriptive shorthand rather than as a substantive claim about behavioral types. Throughout this study, modality refers to a distinct behavioral channel through which accounts participate on the platform: trading, commenting, or reacting. The seven-class typology that follows (Trader Only, Trader+Commenter, etc.) labels wallets by observed presence in each channel and does not imply equivalent investments of effort or attention across channels; we treat the typology as a descriptive partition of accounts rather than a substantive claim about cross-modal participation, and revisit it in both the threshold sensitivity reclassification (RQ1) and the within-event analysis (RQ5). Modalities are not platform affordances per se but observable categories of account action; we use the term to organize our analysis of how accounts distribute their engagement across behaviorally distinct channels.

Data and Methods

Data

We collected data from Polymarket's public-facing API on election-related betting activity between May 14, 2023, and September 5, 2025. These data include betting events, individual markets belonging to events, trades, user-generated comments, and reactions. The data collection spanned 846 days and captured 778,634 unique accounts (defined as unique cryptographic wallet addresses). Accounts engage through three distinct behavioral modalities: (1) trading (placing financial bets on electoral outcomes); (2) commenting (producing political discourse); and (3) reacting (providing social feedback on others' comments). Note that during our observation period, Polymarket trading was not legally available to U.S.-based users under CFTC regulations, though users could access the platform via virtual private networks (VPNs). We cannot determine users' geographic locations solely from wallet addresses, and our analysis focuses on observable platform behaviors rather than user demographics or regulatory compliance. While the data collection period began in May 2023, the vast majority of account activity (99.9% of accounts) occurred on or after October 1, 2023, coinciding with intensified election-related market activity leading up to the 2024 U.S. presidential election and more general platform dissemination and adoption. Figure 1 illustrates the relationships among these data structures and how account activities are linked across trading, commenting, and reacting via proxy wallet identifiers.

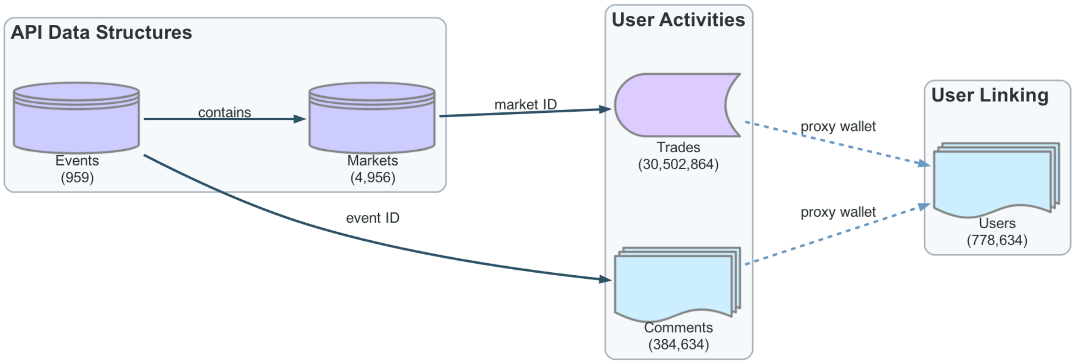

Figure 1. Data collection flow diagram showing the relationships among API data structures (Events and Markets), account activities (Trades and Comments), and account linking. Solid arrows represent direct hierarchical or identifier relationships; dashed arrows represent account linkage via proxy wallet addresses.

Within Polymarket's taxonomy, "events" can consist of a single betting market or a set of related markets, each of which allows accounts to buy and sell shares ("contracts") corresponding to one of two possible outcomes. For example, a Polymarket event that allows accounts to bet on the outcome of New York's 2025 mayoral election will have a market for each candidate, and accounts can purchase or sell "Yes" or "No" contracts in any of these individual markets to place bets "for" or "against" the corresponding candidate's winning the election (e.g., four Zohran Mamdani "yes" contracts). Some events have only one corresponding market, for instance, if only two candidates are facing off in a race (e.g., Who will win the Democratic Nomination in the TX-28 House election: Jessica Cisneros or Henry Cuellar). In these markets, the binary outcome set means that accounts buy or sell contracts corresponding to either candidate (e.g., "Cisneros" or "Cuellar"). Each event has a comments section where accounts can post and reply to one another's posts.

Events. We fetched all current and historical events using Polymarket's "Gamma" API, collecting them in batches of 500. There was no single method we could use to reliably identify election-related events across the entire time period, so we used a multi-pronged approach to narrow the full list to election-related events. First, we searched text fields, including event titles and descriptions, for words like "election" and "primary." Second, we leveraged the optional "category" field in some of the more recent events, which included values (e.g., "Global Politics") that contained election-related event topics. This process identified 959 election-related events, which we saved in raw form. These 959 events serve as our event universe; activity counts are restricted to our observation window (May 14, 2023–September 5, 2025). Within this window, 751 events received at least one comment and 860 recorded at least one executed trade; 749 events saw both, 111 had trades but no comments, and 2 had comments but no trades, yielding 862 "active events" (comments or trades). Many of these events centered around predicting outcomes of specific races (e.g., "Presidential Election Winner 2024," "2024 Democratic Presidential Nominee," "Who will win the 2024 Senate race in Ohio?"), but others are concerned with matters tangential to specific elections (e.g., "Turnout in NYC Mayoral Dem Primary?" "Will Joe Biden drop out of the presidential race?"). All 959 events were manually reviewed to confirm their election relevance.

Markets. All event objects returned by the Gamma API's `/events` endpoint included a list of full data objects for all constituent markets, which eliminated the need to make additional requests to the markets endpoint for the purpose of gathering trade data. After reducing the events dataset from all events to just election-related events, we extracted the markets nested within each event object and saved them separately. The 959 events comprised 4,956 individual markets, with some events aggregating multiple related betting markets (e.g., separate markets for each candidate in a multi-candidate race). Of these 4,956 markets, 4,953 had queryable blockchain identifiers (`conditionId` values) that enabled complete trade data collection via the subgraph (described below); the remaining 3 markets lacked these identifiers and could not be queried. During the observation window, 3,481 of these markets recorded at least one executed trade; the remainder are inactive within the window or fall outside it. Market data enabled us to analyze trades with greater granularity than we could from the relationship between account trades and events.

Trades. To acquire account trades, we queried Polymarket's public subgraph—a GraphQL interface to the platform's on-chain transaction data hosted on The Graph network. We iterated through all election-related markets (identified by their unique `conditionId` values) and collected all order-filled events within our study period (May 14, 2023, through September 5, 2025). For each market, the subgraph query retrieved both maker and taker sides of all transactions, which we then deduplicated by transaction hash to construct complete trade records. Each trade record included the trader's wallet address (which we used to link traders to commenters), market identifier, outcome index (0 or 1, corresponding to outcomes such as "Yes" or "No"), transaction price, trade size, and timestamp. This approach resulted in 30,502,864 trades across all election markets. Dollar volumes reported throughout the manuscript reflect stablecoin-denominated values (USDC) recorded at the time of trade, which maintain a 1:1 peg to USD; we do not apply time-of-trade currency conversions, as all transactions occurred in USD-pegged stablecoins on the Polygon blockchain.

Comments. We again used the Gamma API to retrieve all comments for each election-related event, passing the event ID as a parameter. We iterated over the list of 959 event IDs and made batch requests for each event's comments. The Polymarket comments system allows both top-level event comments and nested replies. Our collection captured the full comment tree structure, preserving parent-child relationships for subsequent analysis of conversation dynamics. Through this process, we collected both top-level comments on events and replies to comments, resulting in a total of 384,586 comments in 751 events. All comments returned by the API included account profile data for the comment's author, including identifiers we used to connect comment authors to account trades. We identified 30,477 unique commenters among the 384,586 comments.

Reactions. Each comment object obtained from the Gamma API `/comments` endpoint included a list of "reactions" (i.e., a "heart") to the comment, accompanied by identifiers for the account responsible for the reaction. We captured 571,523 reactions made by 28,957 unique reactors.

The [ANONYMOUS] IRB reviewed this study (Submission ID [ANONYMOUS]) and determined that the proposed activity is not research involving human subjects.

Account Classification and Behavioral Modalities

We classify accounts based on three behavioral *modalities*: distinct types of platform engagement that accounts can perform. The three modalities are: (1) trading, which involves placing bets on electoral outcomes by buying or selling outcome contracts; (2) commenting, which involves producing political discourse by posting top-level comments or replies on event pages; and (3) reacting, which involves providing social feedback by giving reactions (i.e., a "heart") to others' comments. Based on which modalities each account engaged in during the observation period, we assign accounts to one of seven mutually exclusive account types: (1) Trader Only (accounts that only traded); (2) Commenter Only (accounts that only commented); (3) Reactor Only (accounts that only gave reactions); (4) Trader+Commenter (accounts that both traded and commented but did not react); (5) Trader+Reactor (accounts that both traded and reacted but did not comment); (6) Commenter+Reactor (accounts that both commented and reacted but did not trade); and (7) Trader+Commenter+Reactor (accounts that engaged in all three modalities). This typology is a descriptive partition of accounts by the channels in which they were observed and does not imply equivalent investments of effort or attention across channels. Accounts that engaged in only one modality are classified as single-mode participants, while accounts that engaged in two or more modalities are classified as multi-modal participants. This taxonomy yields 746,458 single-mode accounts (95.9%) and 32,176 multi-modal accounts (4.1%). We define total activity for each account as the sum of their trades, comments, and reactions across the full observation period (May 14, 2023 – September 5, 2025; 846 days).

Account entry is defined as the date of an account's first observed activity (earliest among first trade, first comment, or first reaction). All account-level analyses use the full 846-day observation period, while temporal trend visualizations focus on the period of intensified electoral activity (October 2, 2023 – September 1, 2025; 101 weeks) to better capture patterns around the 2024 U.S. presidential election. Wallet addresses serve as proxies for individual accounts. We cannot determine whether multiple addresses belong to a single person or whether a single address is shared. We do not attempt to identify or exclude bot accounts or algorithmic traders, as there is no reliable ground truth for such classification in pseudonymous blockchain environments.

Methods

To analyze these data descriptively, we combine distributional comparisons, concentration metrics, temporal summaries, and behavioral associations to characterize account engagement across the three modalities. We treat statistical tests as descriptive diagnostics of distributional differences rather than as tests of causal mechanisms or latent user types.

To compare activity levels across account types, we conducted omnibus one-way ANOVA tests ($\alpha = 0.05$) for trading volume (RQ2 Panel A), comment production (RQ3 Panel A), and reaction activity (RQ4 Panel A), with account modality as the group factor representing all seven account types (Trader Only, Commenter Only, Reactor Only, Trader+Commenter, Trader+Reactor, Commenter+Reactor, Trader+Commenter+Reactor). Given the pronounced right skew in all three outcome distributions, we fit log-transformed variants as robustness checks. Nonetheless, the Results report the original-scale omnibus statistics. For the modality count distribution itself (RQ1 Panel A), the analogous question is whether accounts are distributed evenly across the seven modalities rather than whether per-account activity differs in level; we therefore used a chi-square goodness-of-fit test against equal shares ($\alpha = 0.05$) for that panel.

To quantify concentration of activity within and across account types, we calculated Gini coefficients to measure inequality in activity distributions (where 0 indicates perfect equality and 1 indicates maximum concentration), computed top shares showing the percentage of total activity generated by the top 10% and top 1% of accounts, and constructed complementary cumulative distribution functions (CCDFs) showing the proportion of accounts (y-axis, log scale) who performed at least x activities (x-axis, log scale) separately by account modality. To test whether activity distributions differed significantly across account modalities, we applied the Kruskal-Wallis test, a non-parametric alternative to ANOVA suitable for highly skewed distributions, at $\alpha = 0.05$, for per-account activity counts (RQ1 Panel C, RQ2 Panel B, RQ3 Panel B, RQ4 Panel B) and for first-active-week timing across modalities (RQ1 Panel D). CCDFs appear in RQ1 Panel C, RQ2 Panel B, RQ3 Panel B, and RQ4 Panel B. To test whether activity intensity (binned into low/medium/high categories) is independent of account type (RQ1 Panel B), we applied a chi-square test of independence at $\alpha = 0.05$ and report Cramer's V as the effect-size companion.

To assess the sensitivity of these concentration metrics to low-activity accounts, we reestimated each per-modality statistic (account share, Gini, top-10% share, top-1% share) in a sensitivity analysis across $k \in \{1, 2, 5, 10, 25, 100\}$ trades per account. We report the $k = 1$ value (the full population of trading accounts) as the primary analysis level and interleave higher k values at the specific findings they most illuminate; the full sensitivity analysis appears in the appendix.

To examine temporal dynamics, we aggregated account activity by week using Monday-starting weeks over the intensified electoral activity period (October 2, 2023 – September 1, 2025; 101 weeks). We present three types of temporal visualizations: weekly active account counts showing the number of distinct accounts active in each week (RQ1 Panel D, RQ3 Panel D), temporal composition showing the proportion of weekly trading volume or comment activity attributable to each account modality as stacked area plots (RQ2 Panel C, RQ3 Panel D), and cumulative entry patterns tracking the cumulative count of accounts making their first activity in each week stratified by modality (RQ1 Panel D). Key political events (Biden withdrawal, Election Day, and Inauguration) are marked on relevant temporal visualizations to contextualize activity spikes.

For reactions (RQ4 Panel C), per-capita rates are computed as the weekly number of reactions divided by the number of accounts that gave at least one reaction in that week within each account type. Because the active-reactor denominator varies by week (and is smaller than that for trading or commenting), week-to-week fluctuations can appear more pronounced relative to absolute counts.

To examine associations between trading and commenting behaviors among accounts that engaged in both modalities (RQ3 Panel C), we computed Pearson's correlation between log-transformed trade counts and log-transformed comment counts among accounts classified as Trader+Commenter or Trader+Commenter+Reactor (N = 20,203), then fit an analysis of covariance (ANCOVA) model with log-comments as the dependent variable, log-trades as a continuous covariate, and modality as a categorical factor (2 levels: Trader+Commenter, Trader+Commenter+Reactor) to test whether the modality $\times$ log-trades interaction term was significant at $\alpha = 0.05$, indicating that account type moderates the trading-commenting relationship. We visualized these associations using log-log scatterplots with separate regression lines for each modality to show differences in slope and intercept.

To analyze cross-modal participation within individual events (RQ5), we constructed a (wallet, event) panel of 2,364,908 wallet-event pairs by aggregating trades, comments, and reactions to the (wallet, event) level. The panel covers 778,634 wallets and 862 events; we restrict within-event analyses to events that host at least one comment, since events with no commentary cannot exhibit within-event trade-comment co-occurrence by construction. Trades are joined to events via the trade-event identifier. Each (wallet, event) row is assigned a within-event modality using the same seven-class taxonomy used at the wallet level in RQ1, with the wallet's full-period activity replaced by the wallet's activity within that event, so a wallet labeled Trader+Commenter at the wallet level can still be Trader Only in most of their individual events. For Panel B, we compute the signed log-time difference between each wallet's first comment and first trade in an event as $\text{sign}(\Delta) \cdot \log_{10}(|\Delta| + 1)$, where Δ is the difference in hours, restricted to the 31,349 (wallet, event) pairs where the wallet both traded and commented within the event. For Panel C, we compute the per-event Spearman correlation between trade volume and comment count among joint-active wallets within each event, restricted to events with at least 30 joint-active wallets to ensure stable per-event estimates, and report a sensitivity analysis at thresholds of 10, 25, 30, 50, and 100 joint wallets per event. We treat Spearman as the primary per-event coefficient because trade volume spans roughly six orders of magnitude (from single dollars to millions), and per-event whales dominate raw-scale Pearson; for transparency, we also report raw Pearson and log10-volume Pearson alongside Spearman in the per-event correlation distribution. The wallet-aggregate Pearson log-log r from RQ3 Panel C is retained for parallel comparison to the wallet-aggregate trading-commenting association. For Panel D, we classify each wallet active in at least two joint events by whether their per-event ordering is always trade-first, always comment-first, or mixed, and report this consistency across the cell-size sensitivity analysis at k = 2, 5, 10, 25, 100 joint events per wallet.

In summary, RQ1 employs chi-square goodness-of-fit against equal shares for the modality count distribution (Panel A), chi-square tests of independence with Cramer's V on the modality-by-intensity cross-tabulation (Panel B), Gini and CCDF concentration analysis with a Kruskal-Wallis omnibus on per-account activity counts (Panel C), and Kruskal-Wallis tests of first-active-week timing across modalities (Panel D). RQ2 uses ANOVA of trading volume by account type (Panel A), Gini and CCDF concentration analysis of dollar volume (Panel B), and temporal composition of weekly trading by modality (Panel C). RQ3 uses ANOVA of comment counts (Panel A), Gini and CCDF concentration of comments (Panel B), Pearson correlation and ANCOVA for trading-commenting associations (Panel C), and weekly active commenter counts (Panel D). RQ4 uses ANOVA of reaction counts (Panel A), Gini and CCDF concentration of reactions (Panel B), and per-capita temporal reaction rates (Panel C). RQ5 uses within-event modality composition over the (wallet, event) panel (Panel A), signed log-time differences between first trade and first comment within events (Panel B), per-event Spearman correlations of trade volume against comment count with raw and log Pearson companions (Panel C). Per-wallet ordering consistency across joint events is reported across the joint-event-cutoff sensitivity analysis (Panel D). Formal time-series modeling (e.g., seasonality decomposition, autocorrelation analysis) represents a natural extension of the temporal patterns we document but exceeds the descriptive scope of this study.

Results

We analyzed 778,634 unique accounts that participated in election betting markets over the study period. Accounts engaged across three behavioral modes: trading (placing bets on electoral outcomes), commenting (producing political discourse), and reacting (providing social feedback on others' comments). The analysis proceeds from foundational account taxonomies (RQ1) through behavioral patterns in financial engagement (RQ2), communication (RQ3), and social interaction (RQ4).

RQ1: Account Characteristics and Cross-Mode Participation

This research question documents who participates in election betting markets and how accounts distribute engagement across behavioral modes (see Figure 2).

Modality Distribution (Panel A). First, we find that the platform is overwhelmingly populated by accounts whose only observed activity is trading (Figure 2A). Trader Only comprises 94.4% (N = 734,867), with Commenter Only (0.7%, N = 5,470) and Reactor Only (0.8%, N = 6,121) accounting for the remaining single-mode accounts. The aggregate 95.9% single-mode share thus largely restates the prevalence of trading-only participation rather than indexing a substantive split between specialist and multi-modal engagement; we treat the seven-class typology as a descriptive partition rather than a substantive claim about behavioral types. Multi-modal participation accounts for 4.1%—chiefly Trader+Commenter+Reactor (1.4%, N = 10,863), followed by Trader+Commenter (1.2%, N = 9,340), Trader+Reactor (0.9%, N = 7,169), and Commenter+Reactor (0.6%, N = 4,804). The distribution differs markedly from uniform allocation ($\chi^2(6) = 4{,}079{,}407.47$, $p < .001$). Under a sensitivity analysis in which a wallet must perform at least k events in every modality counted toward its label, the multi-mode share declines to 1.8% at $k = 10$ and 1.7% at $k = 100$. Per-cohort counts at each level appear in Appendix A.1. Within each modality, low-activity participation is substantial: 12.6% of trading accounts execute exactly one trade, 41.1% of commenting accounts post exactly one comment, and 34.3% of reacting accounts leave exactly one reaction; the full per-modality sensitivity analysis appears in Appendix A.2.

Engagement Intensity (Panel B). Next, we find that engagement intensity stratifies systematically by account type (Figure 2B). We bin accounts on absolute activity counts (1, 2–4, 5–9, 10–24, 25–99, 100+ events per account) so the bins are mutually exclusive between sensitivity-analysis levels and remain interpretable across modalities and stable under sample changes; the full per-modality contingency appears in Appendix A.3. Account type and activity intensity are significantly associated under these bins ($\chi^2(30) = 45{,}472.59$, $p < .001$, Cramér's V = 0.108), suggesting distinct engagement profiles across account categories. Multi-modal accounts, particularly those in the Trader+Commenter+Reactor category, concentrate in higher activity bins (27.8% sit in the 100+ bin). In comparison, single-mode Commenter Only and Reactor Only accounts cluster in the lowest bins (59.6% and 48.7%, respectively, sit in the singleton bin). This stratification demonstrates that the breadth of participation modes correlates with the depth of engagement, underscoring the need to quantify the inequality structure underlying these differences in intensity.

Activity Concentration (Panel C). Furthermore, we observe high concentration in overall activity that varies systematically across account types (Figure 2C). The complementary cumulative distribution function (CCDF) reveals a Gini coefficient of 0.850 for total activities: the top 10% of accounts account for 78.7% of all activities, and the top 1% account for 57.2%. Concentration is not driven entirely by low-activity accounts: among accounts with at least 10 total events, the Gini coefficient is 0.769, and among the 4.5% of accounts with at least 100 events, it is 0.759. The full pooled sensitivity analysis (k = 1, 2, 5, 10, 25, 100) appears in Appendix A.2. Activity distributions differ significantly across account modalities (Kruskal–Wallis $\chi^2(6) = 30{,}050.21$, $p < .001$). Trader+Reactor accounts exhibit the highest inequality (Gini = 0.947), while Commenter Only accounts show the lowest (Gini = 0.501);

Trader+Commenter+Reactor accounts fall between these extremes at Gini = 0.918, with Reactor Only at 0.885 and Trader Only at 0.809. The Trader+Reactor / Commenter Only contrast persists across the sensitivity analysis: Trader+Reactor remains the most-concentrated cohort (Gini = 0.929 at $k = 10$, 0.861 at $k = 100$) while Commenter Only remains the least-concentrated (Gini = 0.422 at $k = 10$, 0.085 at $k = 100$, n = 6); the full per-modality sensitivity analysis appears in Appendix A.2. These patterns indicate that behavioral specialization is associated with heavy-tailed distributions of activity levels.

Temporal Entry Patterns (Panel D). Finally, entry timing varies systematically across account types, complementing the cross-sectional composition and inequality patterns (Figure 2D). Account types entered at significantly different times (Kruskal–Wallis $H(6) =$ 2,665.29, $p < .001$), revealing temporal stratification in adoption patterns. Trader-only accounts dominated early platform adoption, entering steadily from the beginning of the observation period. Multi-modal accounts, particularly those engaging in all three modes and Trader+Commenter, entered the platform later in its lifecycle. This temporal differentiation completes the characterization of account heterogeneity by adding a temporal dimension to the compositional, intensity, and inequality patterns documented in Panels A-C.

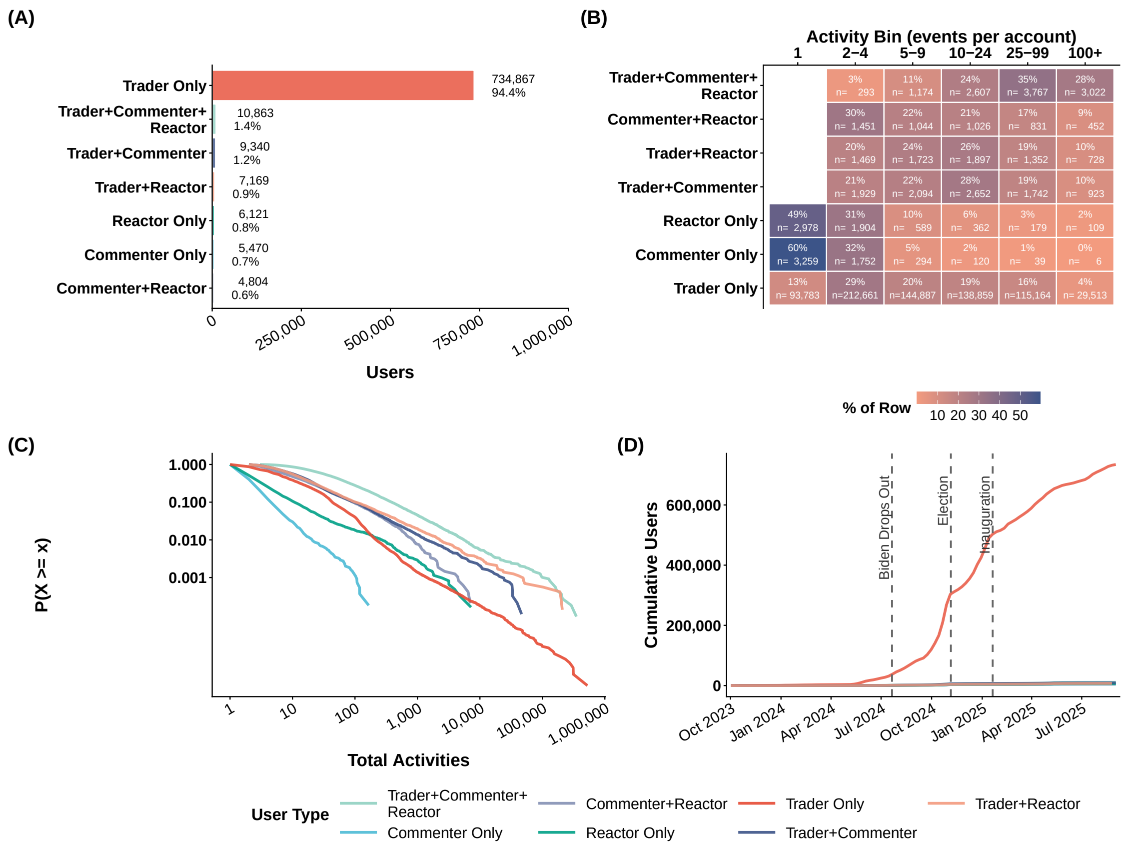

Figure 2. Account Composition and Cross-Mode Participation. (A) Distribution of 778,634 accounts across seven participation modalities, showing dominance of single-mode specialization (95.9%). (B) Cross-tabulation of account modality by absolute activity bins (1, 2–4, 5–9, 10–24, 25–99, 100+ events per account), revealing systematic stratification ($\chi^2(30) = 45{,}472.59$, $p < .001$, Cramér's V $= 0.108$). Cells report row percentages (each row sums to 100%), so each cohort's intensity profile can be read directly; the column-percentage view is reported in Appendix Table 6. (C) Complementary cumulative distribution functions (CCDFs) of total activity by account type, demonstrating high concentration (Gini = 0.850) with variation across modalities (Kruskal–Wallis $\chi^2(6) = 30{,}050.21$, $p < .001$). (D) Cumulative account entry over time by modality, showing temporal stratification in platform adoption (Kruskal–Wallis $H(6) = 2{,}665.29$, $p < .001$).

RQ2: Financial Engagement in Political Markets

Building on the RQ1 taxonomy, we examine trading behaviors across account types. Figure 3 progresses from cross-sectional distributional differences (Panel A) to concentration structure (Panel B), then to temporal composition (Panel C).

Trading Volume by Account Type (Panel A). Trading volume is heavy-tailed across the wallet population: most traders place small-volume bets while a long right tail accounts for the bulk of dollar flow (Figure 3A). Disaggregating by account type, Trader+Commenter+Reactor accounts exhibit the highest median trading volumes, followed by Trader+Commenter, Trader+Reactor, and Trader Only accounts ($F(3, 762235) = 647.02$, $p < .001$). Violin distributions indicate substantial within-type heterogeneity with long right tails, reflecting small subsets of high-volume traders in each category. This monotonic ordering is largely a mechanical consequence of the fact that wallets active in more modalities tend to be more active overall; the within-event coupling question itself is taken up in RQ5. These cross-sectional differences warrant examination of how trading activity concentrates within and across account types.

Trading Concentration (Panel B). Trading volume is sharply concentrated (Figure 3B). Overall inequality reaches Gini = 0.951: the top 10% of traders account for 92.1% of total dollar volume, and the top 1% for 73.7%. Volume concentration persists when low-activity traders are excluded: among accounts with at least 10 trades, the volume Gini is 0.924, and among the 4.3% of accounts with at least 100 trades, it is 0.930. The top 1% share declines monotonically across the sensitivity analysis from 73.7% at $k = 1$ to 64.0% at $k = 100$, indicating that the singleton mass is not what props up tail concentration. The full pooled trade-volume sensitivity analysis appears in Appendix A.2. A view of the low-activity tail (Appendix A.4) makes the trade singleton mass directly visible without the log-transformation used for the CCDF. This concentration level (Gini = 0.951) substantially exceeds the overall activity concentration from RQ1 (Gini = 0.850), indicating that financial stakes concentrate more sharply than overall behavioral participation. Disaggregating by account type, modality profiles diverge in rank distribution (Kruskal–Wallis $\chi^2(3) = 9{,}904.63$, $p < .001$), tracking the same mechanical pattern noted for Panel A; this rank divergence is parallel structure rather than direct evidence of trading-discourse coupling.

Temporal Trading Composition (Panel C). Next, turning to the temporal dimension, weekly trading composition shows systematic variation in account type participation across the electoral cycle (Figure 3C). During the early period (October 2023 - March 2024), Trader Only accounts constitute a relatively small proportion of trading activity, while multi-modal accounts (particularly Trader+Reactor and Trader+Commenter+Reactor) account for substantial shares. Following Biden's withdrawal (July 2024), composition shifts toward greater Trader Only dominance. The general election period (November 2024) marks the peak Trader Only proportional contribution, reaching approximately 75-80% of weekly trading volume.

Post-election, the composition shows increased volatility, with multi-modal accounts regaining a proportional share, particularly around the inauguration period (January 2025). Throughout the observation window, Trader Only accounts consistently provide the majority of trading volume, but their proportional dominance varies substantially—lowest in early 2024 and highest around the general election. This compositional view reveals that different account types contribute differentially across the electoral cycle, with single-mode traders dominating high-stakes moments while multi-modal engagement shows greater relative presence during lower-intensity periods.

RQ3: Communication Behaviors in Political Discourse

Next, we examine commenting activity (see Figure 4).

Comment Production by Account Type (Panel A). Comment production is heavy-tailed across accounts active in commenting: most commenters produce a small number of comments, while a long right tail accounts for the bulk of discourse (Figure 4A). Disaggregating by account type, Trader+Commenter+Reactor accounts produce the most comments on average, followed by Commenter+Reactor, Trader+Commenter, and Commenter Only accounts ($F(3, 30473) = 257.80$, $p < .001$). Distributions are positively skewed within each type, consistent with the presence of a minority of prolific commenters in each group. The same mechanical relationship that produced the trading ordering in RQ2 is at work here, since wallets present in more modalities also tend to accumulate more comments overall; we report the contrast for

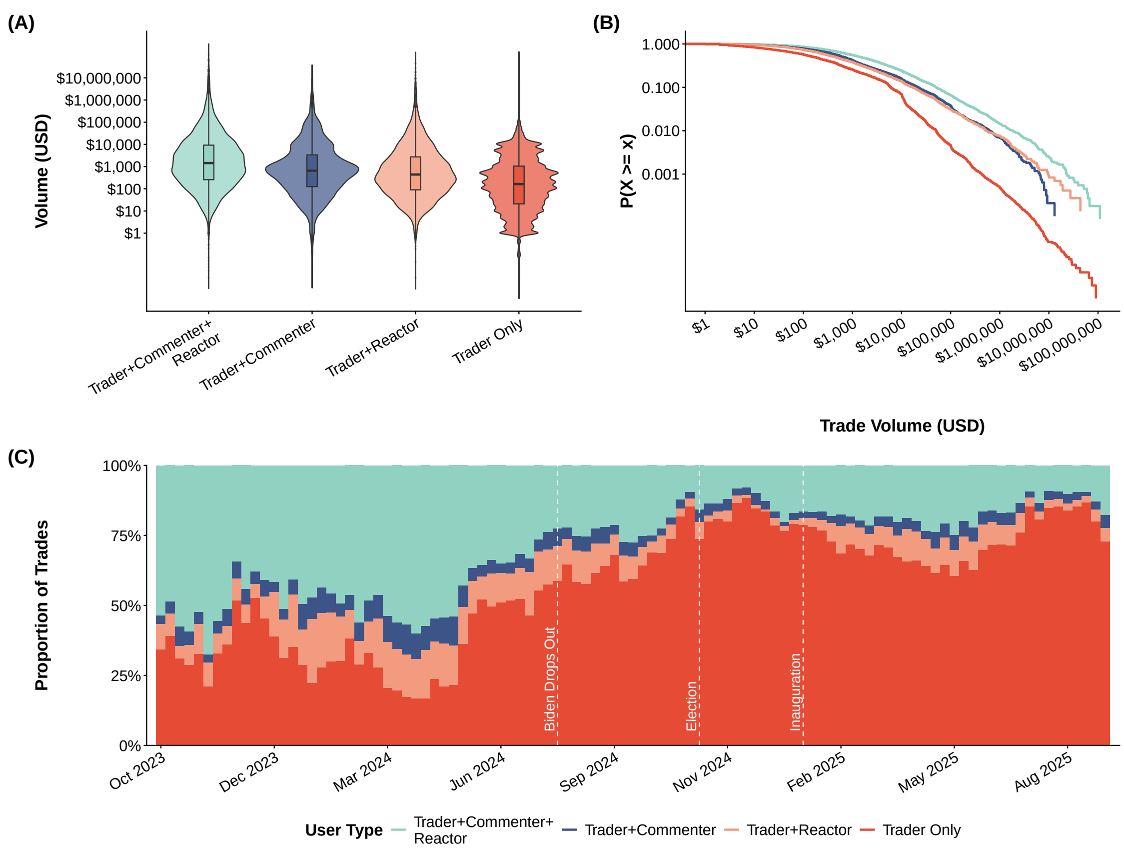

Figure 3. Financial Engagement in Political Markets. (A) Violin and box plots of trading volume by account type, showing significant differences ($F(3, 762235) = 647.02$, $p < .001$) with Trader+Commenter+Reactor accounts exhibiting the highest median volumes. (B) CCDF of dollar volume by modality, revealing extreme concentration (Gini = 0.951; Kruskal–Wallis $\chi^2(3) = 9{,}904.63$, $p < .001$) with top 10% accounting for 92.1% of total volume. (C) Weekly trading composition from October 2023 through August 2025 (101 weeks), showing temporal clustering around high-salience political events (Biden withdrawal, election, inauguration).

completeness and rely on the within-event panel in RQ5 for the cross-modal coupling question itself. These differences position us to directly quantify concentration.

Comment Concentration (Panel B). Comment production is highly concentrated (Figure 4B). Overall, the Gini coefficient is 0.814: the top 10% of commenters account for 76.4% of all comments, and the top 1% account for 36.0%. Comment-volume concentration attenuates as low-activity commenters are excluded, but does not collapse: among accounts with at least 10 comments, the comment Gini is 0.625, and among the 2.3% of commenters with at least 100 comments, it is 0.403. The full pooled comment sensitivity analysis appears in Appendix A.2. A two-group decomposition of the comment Gini into singleton-vs-active components shows that 46.5% of overall comment inequality comes from the gap between singleton commenters and the rest, and 53.5% from heterogeneity among accounts with at least two comments; the full per-modality decomposition appears in Appendix A.5. A view of the low-activity tail (Appendix A.4) makes the singleton mass directly visible without the log-transformation used for the CCDF. This concentration level (Gini = 0.814) is lower than both overall activity concentration from RQ1 (Gini = 0.850) and trading volume concentration from RQ2 (Gini = 0.951), positioning discursive engagement as the least concentrated of the measured behaviors. Disaggregating by account type, concentration patterns vary across modality profiles (Kruskal–Wallis $\chi^2(3) = 6{,}501.26$, $p < .001$), echoing the parallel structure noted for Panel A rather than carrying its own cross-modal interpretation.

Trading-Commenting Association (Panel C). Among accounts active in both behaviors, log-transformed trading and commenting frequencies are positively associated (Pearson $r = 0.28$, 95% CI [0.27, 0.29], $n = 20{,}203$, $p < .001$; Figure 4C). The wallet-aggregate association attenuates but does not collapse under joint thresholding. Requiring at least 10 trades and at least 10 comments per wallet yields Pearson log-log $r = 0.18$, 95% CI [0.15, 0.22] (n = 2,915), and at the top level ($k = 100$ of each) $r = 0.16$, 95% CI [0.02, 0.29] ($n = 200$); the Pearson CI excludes zero at every level from $k = 1$ to $k = 100$. The Spearman rank correlation tracks this trajectory as a robustness companion (Spearman $r = 0.24$ at $k = 1$ and 0.14 at $k = 10$), with its CI excluding zero through $k = 25$ but including zero at $k = 100$ where the joint-active cohort shrinks to 200 wallets. The full joint-threshold table appears in Appendix A.6. Correlations vary by account type: Trader+Commenter+Reactor accounts show $r = 0.28$ (95% CI [0.26, 0.29], n = 10,863), while Trader+Commenter accounts exhibit $r = 0.15$ (95% CI [0.13, 0.17], n = 9,340). Account modality significantly moderates the trade-comment relationship (ANCOVA, $F(2, 20199)$ = 2,628.86, $p < .001$), indicating heterogeneity across account types in this relationship. The moderate association is consistent with only partial coupling between financial and discursive engagement. RQ5 directly examines the within-event analog of this association. Once both behaviors are localized to the same event, the per-event Spearman distribution brackets rather than exceeds this wallet-aggregate value, and only 21.8% of wallet-event combinations in the combined cohort realize both behaviors within the same event (see RQ5).

Temporal Comment Dynamics (Panel D). Commenting around major political events, all four groups exhibit event-aligned spikes that vary in absolute magnitude (Figure 4D). Two pronounced spikes appear in the series. The U.S. general election week (week of

November 4, 2024) reached 53,156 weekly comments across the four commenting modalities, with Trader+Commenter+Reactor (teal) contributing 25,668 and Commenter+Reactor (purple) contributing 19,732, reaching comparable absolute magnitudes. A second spike of similar scale occurs in the week of May 12, 2025, totaling 48,480 weekly comments and concentrated within the Trader+Commenter+Reactor group (37,722 that week); this peak coincides with the Romanian presidential runoff (May 18, 2025), a high-stakes rerun ordered after the Romanian Constitutional Court annulled the November 2024 first-round results in early December 2024 based on declassified intelligence alleging Russian-backed influence operations promoting far-right candidate Călin Georgescu across social platforms. Biden's withdrawal from the 2024 U.S. presidential race in July 2024 produced only modest commenting activity by comparison. Commenter Only (light blue) and Trader+Commenter (dark blue) remain numerically smaller throughout the series but also rise sharply during peak weeks (Commenter Only reaching 3,034 and Trader+Commenter reaching 4,722 during the November 2024 election week). These dynamics reveal that commenting concentrates temporally around specific political moments, with cross-group differences manifesting in absolute magnitude rather than the presence of event-driven sensitivity.

RQ4: Social Reaction and Community Engagement

Extending from discourse, we analyze reaction-giving across account types (see Figure 5).

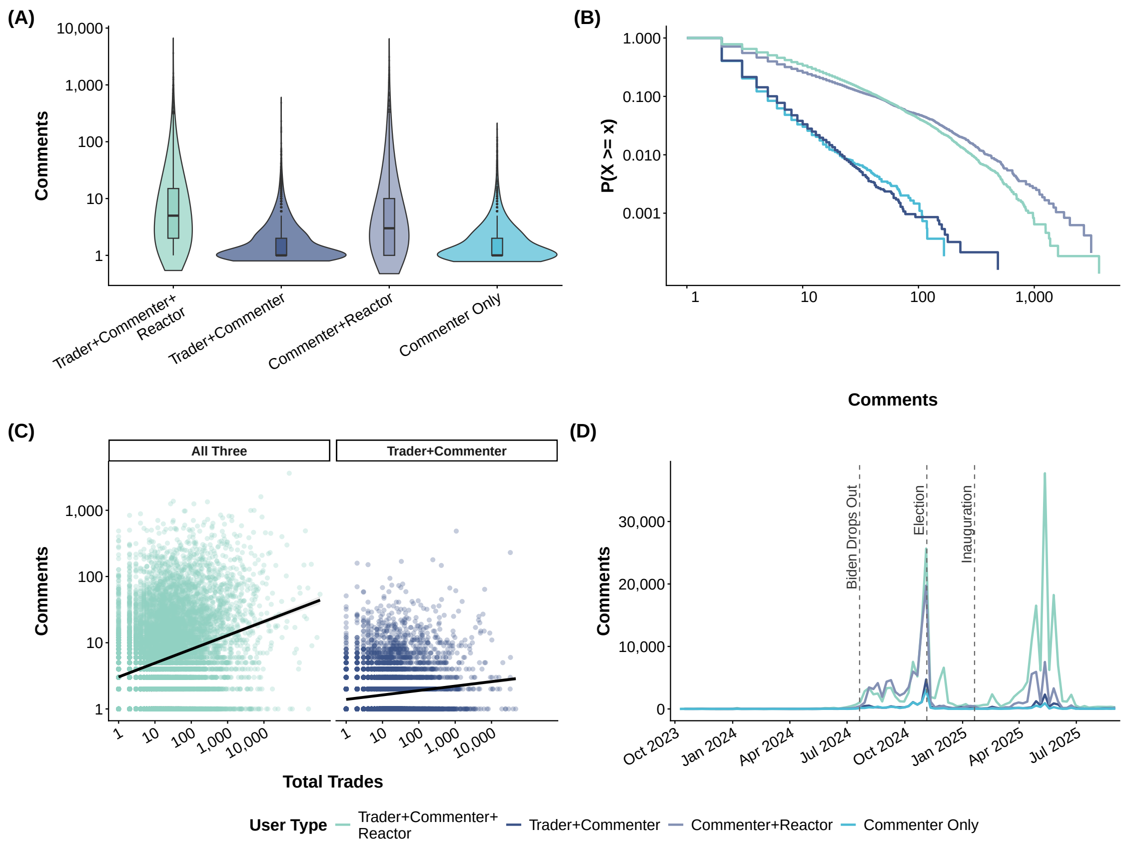

Figure 4. Communication Behaviors in Political Discourse. (A) Violin and box plots of comment production by account type ($F(3, 30{,}473) = 257.80$, $p < .001$). (B) CCDF of comments by modality showing substantial concentration (Gini = 0.814; Kruskal–Wallis $\chi^2(3) = 6{,}501.26$, $p < .001$). (C) Log-log scatterplot of trading vs. commenting frequencies for Trader+Commenter ($r = 0.15$, 95% CI [0.13, 0.17]) and Trader+Commenter+Reactor ($r = 0.28$, 95% CI [0.26, 0.29]) accounts, with separate regression lines. Modality significantly moderates the association (ANCOVA $F(2, 20{,}199) = 2{,}628.86$, $p < .001$). (D) Weekly comment counts by account type, with dashed reference lines marking Biden's July 2024 withdrawal from the 2024 U.S. presidential race, the November 2024 U.S. general election, and the January 2025 U.S. presidential inauguration.

Reaction Activity by Account Type (Panel A). Reaction activity is heavy-tailed across accounts active in reacting: most reactors give a small number of reactions while a long right tail accounts for the bulk of social feedback (Figure 5A). Disaggregating by account type, Trader+Commenter+Reactor accounts give the most reactions on average, followed by Commenter+Reactor, Trader+Reactor, and Reactor Only accounts ($F(3, 28953) = 43.30$, $p <$ .001). Distributions are positively skewed across groups, as in commenting, with medians below the means, indicating that small subsets of prolific reactors drive aggregate social engagement within each account type. As in RQ2 and RQ3, the monotonic ordering largely reflects the mechanical fact that wallets in more modalities accumulate more activity in each one; the contrast is reported in parallel with the earlier RQs rather than as direct evidence about cross-modal participation. The F-statistic for reactions (43.30) is markedly smaller than those for trading (647.02) and commenting (257.80), indicating that modality breadth differentiates reaction activity less strongly than it differentiates trading or commenting.

Reaction Concentration (Panel B). Attention is concentrated, but its magnitude lies between comment production and trading volume (Figure 5B). Overall, the Gini coefficient is 0.845: the top 10% of reactors account for 79.7% of all reactions, and the top 1% for 42.0%. Reaction concentration attenuates as low-activity reactors are excluded but does not collapse: among accounts with at least 10 reactions the reaction Gini is 0.692, and among the 3.3% of reactors with at least 100 reactions it is 0.495. The full pooled reaction sensitivity analysis appears in Appendix A.2. A two-group decomposition of the reaction Gini into singleton-vs-active components shows that 38.6% of overall reaction inequality comes from the gap between singleton reactors and the rest, and 61.4% from heterogeneity among accounts with at least two reactions; the full per-modality decomposition appears in Appendix A.5. A view of the low-activity tail (Appendix A.4) makes the singleton mass directly visible without the log-transformation used for the CCDF. This concentration level (Gini = 0.845) exceeds comment concentration (Gini = 0.814) but remains well below trading volume concentration (Gini = 0.951), positioning social engagement as highly but not extremely concentrated. Disaggregating by account type, rank distributions differ across modality profiles (Kruskal–Wallis $\chi^2(3) = 3{,}478.37$, $p < .001$), continuing the parallel-structure pattern noted in the earlier RQs without itself indexing cross-modal participation.

Temporal Reaction Patterns (Panel C). Per-capita reaction rates fluctuate with event salience, exhibiting episodic, event-aligned spikes interspersed with low baselines (Figure 5C). Spikes cluster following Biden's withdrawal, around the 2024 U.S. general election, and near the January 2025 inauguration, with the largest excursions among Reactor Only and Commenter+Reactor; Trader+Reactor remains comparatively muted. Trader+Commenter+Reactor accounts sustain higher per-capita rates overall, but also display bursts in mid-2025. These patterns indicate that social engagement responds to major political events, though its magnitude varies across account types.

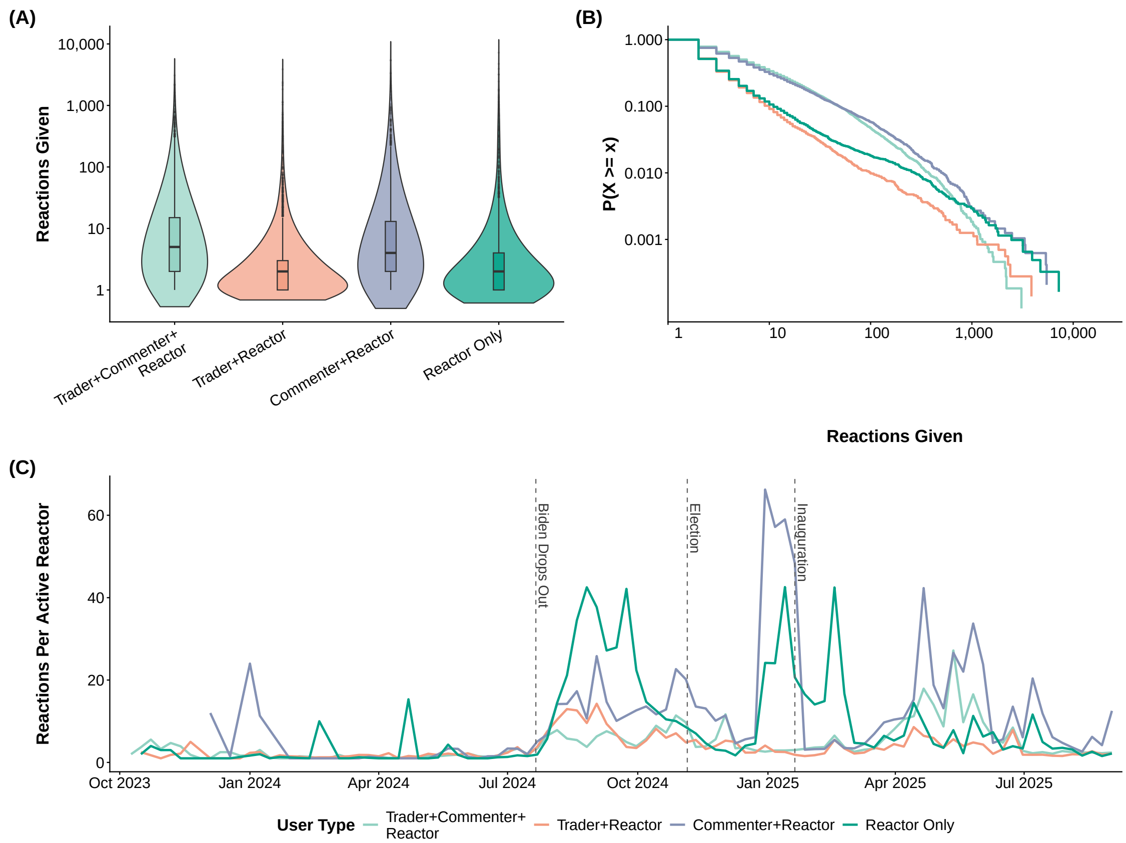

Figure 5. Social Reaction and Community Engagement. (A) Violin and box plots of reaction activity by account type, showing significant but modest differences ($F(3,\ 28{,}953) = 43.30$, $p < .001$). (B) CCDF of reactions by modality, revealing a high concentration (Gini = 0.845; Kruskal–Wallis $\chi^2(3) = 3{,}478.37$, $p < .001$) between commenting and trading levels. (C) Per-capita reaction rates fluctuate with event salience, showing pronounced spikes following Biden's withdrawal, around the 2024 U.S. general election, and near the January 2025 inauguration, with the largest excursions among Reactor Only and Commenter+Reactor. Note: Analyses restricted to accounts with valid wallet linkage.

RQ5: Within-Event Cross-Modal Participation

Extending the cross-sectional taxonomy used in RQ1–RQ4, we analyze cross-modal participation within events using a (wallet, event) panel of 2,364,908 wallet-event pairs across 778,634 wallets and 862 events. Figure 6 progresses from per-wallet within-event joint trade-and-comment realization (Panel A) to within-event temporal ordering (Panel B), then to per-event correlation (Panel C) and cross-event timing consistency (Panel D).

Within-Event Joint Trade-and-Comment Realization (Panel A). The within-event analog of the modest wallet-aggregate trading-commenting correlation reported in RQ3 is the share of touched events in which a wallet that both trades and comments somewhere in the observation period realizes both behaviors within the same event. Among the 20,203 wallets that traded and commented at least once in the observation period, both behaviors co-occur within the same event in only 21.8% of wallet-event combinations (Figure 6A), and 4.7% of these wallets never realize both behaviors within a single event. The realization rate falls further when the cohort is restricted to wallets active in more events, with the full minimum-event-count sensitivity analysis reported in Appendix A.7. Within-event coupling between trading and commenting is therefore weaker than the wallet-aggregate measure alone would suggest. Within-wallet variation in the content of those comments themselves is small: across 2,032 wallets that authored comments in both joint and comment-only events, we find no detectable difference in length, readability, or sentiment between the two strata (Appendix A.8).

Temporal Ordering Within Joint Pairs (Panel B). Among the 31,349 wallet-event pairs in which the same wallet both traded and commented within the event, the trade strictly precedes the first comment in 90.2% of pairs, with a median signed delta of 20.9 hours and 26.8% of pairs occurring within one hour of each other (Figure 6B). The same clustering holds across all 245,198 comments authored in these joint pairs and not just first comments: 31.8% fall within an hour of the wallet's nearest trade in the same event and 71.1% fall within a day (Appendix A.9). The trade-first pattern is also stable within wallets across events: of wallets with at least two joint events ($n = 4{,}768$), 71.9% are always trade-first and only 1.0% are always comment-first; at the more demanding cutoff of five joint events ($n = 797$), 46.5% are always trade-first and 53.5% are mixed (Figure 6D). Comments authored within joint pairs also receive more reactions than comments by wallets that did not trade the same event (Appendix A.10).

Per-Event Trade–Comment Correlation (Panel C). The within-event analog of the wallet-aggregate trading-commenting correlation reported in RQ3 is the distribution of per-event Spearman correlations between trade volume and comment count among joint-active wallets within each event. At a minimum cell size of 30 joint wallets per event, 135 of 704 joint events (19.2%) are estimable, with a median per-event Spearman of 0.247, a wallet-weighted mean of 0.220, and an interquartile range of 0.168 to 0.330 (Figure 6C). The central tendency is stable across the cell-size sensitivity analysis (median Spearman ranges from 0.245 to 0.264 across $k = 10$ to 100, while the share of joint events meeting each cutoff falls from 40.5% to 6.5%); the full sensitivity analysis appears in Appendix A.11. Compared to the RQ3 wallet-aggregate Spearman of 0.243 across the same combined trader-and-commenter cohort, the per-event distribution brackets rather than exceeds the wallet-aggregate value, so within-event coupling between trading volume and comment count is on average no stronger than the cross-sectional measure.

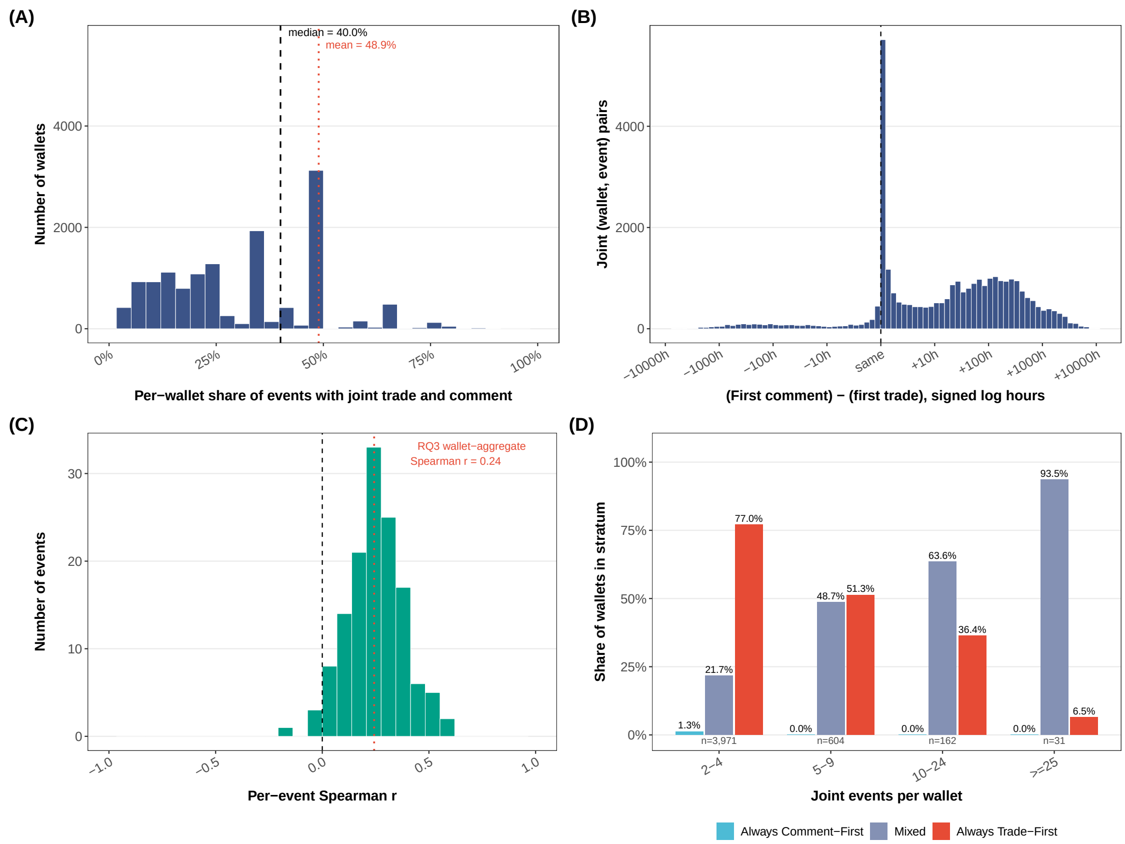

Figure 6. Within-Event Joint Trade-and-Comment Participation. (A) Histogram of the per-wallet share of touched events that host both a trade and a comment by that wallet within the event, among the 20,203 wallets that traded and commented at least once in the observation period (dashed line: per-wallet median; dotted line: per-wallet mean). The realization rate across wallet-event combinations in this cohort is 21.8%. (B) Signed log-time difference between each wallet's first comment and first trade in the same event ($n = 31{,}349$ joint pairs); 90.2% of pairs are trade-first, dashed line at zero. (C) Per-event Spearman correlation between trade volume and comment count among joint-active wallets ($n = 135$ of 704 joint events with ≥ 30 joint wallets); median 0.247, wallet-weighted mean 0.220. Dashed line at zero; dotted line marks the RQ3 wallet-aggregate Spearman ($r = 0.243$). (D) Cross-event timing consistency across the joint-event-cutoff sensitivity analysis at $k = 2$, 5, 10, 25, 100 joint events per wallet, with each wallet classified as always trade-first, always comment-first, or mixed. Stratum sizes (n) are annotated under each cutoff group; no wallets meet the $k = 100$ cutoff. Always-trade-first share moves from 71.9% at $k \geq 2$ to 6.5% at $k \geq 25$ as the share of mixed wallets rises correspondingly. Note: Within-event analyses are restricted to events that host at least one comment.

Discussion

The patterns we find provide initial empirical baselines for wallet-level participation in hybrid political-financial environments. We interpret them through descriptive lenses rather than proposing causal mechanisms. Where we connect findings to existing frameworks, these connections are offered as conjectures to motivate future explanatory research, not as tested theoretical claims. The emphasis is therefore on sharpening what the descriptive evidence can support: visibility makes financial, discursive, and reactive traces jointly observable, but the observed traces remain unevenly distributed and only partially coupled. We make four primary contributions.

First, within-event cross-modal co-occurrence is the exception rather than the rule, and when it does occur, it is temporally ordered. Among the 20,203 wallets that both traded and commented at least once in the observation period, both behaviors co-occur within the same event in only 21.8% of wallet-event combinations (Figure 6A), and among the 31,349 wallet-event pairs where both behaviors did co-occur, 90.2% are trade-first. The per-event Spearman correlation between trade volume and comment count among joint-active wallets has a median of 0.247. Second, participation is predominantly single-mode rather than multi-modal, and that single-mode participation is overwhelmingly financial: 94.4% of accounts trade only, with Commenter Only (0.7%) and Reactor Only (0.8%) comprising small additional single-mode subpopulations. Only 4.1% of accounts participate across two or more modalities. Among the minority who do engage in multiple modes, the correlation between trading volume and commenting volume is modest ($r = 0.28$), indicating partial coupling at best. This pattern is consistent with what platform affordances theory terms differential “action possibilities” (Treem et al., 2020): trading, commenting, and reacting differ in cost, visibility, and effort. Trading requires capital and risk tolerance, commenting requires compositional effort, and reactions require only the investment of attention. Yet the weak coupling between trading and discourse indicates that financial and discursive logics operate semi-independently even when the same accounts engage in both. Some accounts trade extensively with minimal commentary, while others comment prolifically with modest trading, implying distinct participation orientations even under shared visibility. In financial contexts, prior work shows that online discourse can have market-relevant consequences (Antweiler & Frank, 2004) and that prediction settings elicit voluntary knowledge provision alongside financial engagement (Qiu & Kumar, 2017), yet our findings suggest these behaviors remain largely decoupled at the account level. This partial coupling ($r = 0.28$) may reflect distinct participation costs associated with each behavior: trading requires capital deployment and risk tolerance, while commenting requires compositional effort and social exposure. These costs create different participation thresholds that accounts cross independently, such that an account that trades heavily need not comment frequently (and vice versa). The modest correlation shows that financial and discursive engagement remain empirically distinct even among accounts that participate in both modes.

Third, concentration is pronounced across all modes but varies systematically by modality. Trading exhibits the strongest inequality (Gini = 0.951), reactions intermediate concentration (0.845), and discourse the lowest, though still substantial (0.814). Trading, the most consequential behavior with direct financial stakes, exhibits the strongest concentration, consistent with entry barriers related to capital, risk tolerance, and market knowledge, and with strategic silence among less confident participants when positions are publicly observable. Commenting requires compositional effort and discursive engagement, and shows a substantially lower concentration, as discourse serves multiple functions—argumentation, information-seeking, social bonding—that do not uniformly require strategic restraint. Reactions, despite being low-cost actions, remain concentrated, a pattern consistent with strategic social positioning rather than casual feedback in visible communities. The systematic modality gradient thus documents how visibility co-occurs with consequence across participation modes, with higher-stakes behaviors exhibiting stronger inequality. On this platform, higher-consequence observable behaviors coincide with sharper participation inequality. Recent platform-wide evidence suggests that the heavy right tail of trading on Polymarket is plausibly informational: a small subset of accounts (on the order of 3%) appears to drive the majority of price discovery, with their trades predicting both future prices and resolution outcomes (Gómez-Cram et al., 2026). The trading concentration we document may therefore reflect informational stratification alongside participatory inequality. These sensitivity analysis results show that singleton mass and engaged heterogeneity coexist across all modalities rather than substituting for one another. A two-group decomposition of each per-account Gini into singleton-vs-active components attributes 14.5% of trade-count inequality, 46.5% of comment inequality, and 38.6% of reaction inequality to the gap between singleton accounts and the rest, with the remaining 85.5%, 53.5%, and 61.4% reflecting heterogeneity among accounts with at least two events in the modality (Appendix A.5, with cross-cutoff sensitivity in Table 9). The trade > reaction > comment ordering survives the full sensitivity analysis: at $k = 10$ events the per-modality Ginis are 0.924 (trading volume), 0.692 (reactions), and 0.625 (comments), and at $k = 100$ they are 0.930, 0.495, and 0.403 (Appendix A.2).

Fourth, temporal patterns differentiate modalities in substantively meaningful ways. Trading composition shifts markedly across the electoral cycle: single-mode traders account for a substantially higher share of trading volume near elections than in earlier periods—a temporal concentration of trading-only activity around high-stakes electoral moments. In contrast, during lower-intensity periods, participants are more likely to be from accounts that are more prolific across modalities. Commenting exhibits pronounced event responsiveness. Reactions vary more steadily across time, consistent with their role as ambient social feedback and community maintenance rather than event-driven argumentation. Together, these temporal contrasts are consistent with complementary participation patterns: trading as punctuated commitment concentrated at high-stakes moments, discourse as episodic argumentation responsive to diverse political events, and reactions as background social feedback maintaining community presence across periods of varying intensity. As a conjecture for future research, these patterns may reflect observability-salience interactions: trading concentrates when outcome stakes peak, while discourse persists as interpretive work beyond immediate events.

For computer-mediated communication scholarship, these platforms are sites of political talk in which financial stakes create accountability mechanisms absent from traditional social media. When accounts debate an election outcome while holding a visible financial position on it, political expression operates under a form of consequentiality distinct from that of reputation-only platforms. The within-event evidence shows that this consequentiality does not simply suppress discourse: a subset of accounts both trade and comment within the same political moments, with trading almost always preceding commenting, producing political discussion under conditions where their financial commitments are observable to interlocutors. This structure makes betting platforms a theoretically important case for studying political expression under consequence, as the visibility of financial positions may relate to both the content and credibility of political arguments in ways that purely reputational social media environments do not.

Limitations

This descriptive study's validity faces constraints across three dimensions. Internal validity concerns the accuracy of our measurements within this specific context. Analyses depend on wallet linkage, which can reduce sample completeness or linkage if individuals use multiple wallets. Sybil behavior (a single person operating multiple wallet addresses) is prevalent in Web3 environments and could inflate account counts or distribute a single person's activity across multiple accounts. The unit of analysis throughout this paper is the wallet address. To the extent that individuals operate multiple wallets, our wallet-level findings describe the distribution of activity across addresses rather than across persons. Polymarket's design permits this pattern (Douceur, 2002), and the public visibility of holdings, profit-and-loss, and trading histories makes fragmentation a possible response by participants who wish to obscure position sizing, separate strategies, or limit social exposure of losses. Each of the main findings we report could reflect strategic identity fragmentation in addition to underlying behavior: participation counts and the 95.9%/4.1% single-mode/multi-modal partition could be inflated by individuals splitting activity across distinct wallets, and the concentration metrics could be deflated relative to what would obtain at the person level. Recent on-chain methodological work quantifies the magnitude of these behaviors on Polymarket: graph-based detection flags roughly 25% of historical trading volume as wash trading (Sirolly et al., 2025), USDC-flow clustering identifies coordinated wallet groups that capture a disproportionate share of winner profit (Liu, 2026), and a composite anomaly score on Polymarket and Kalshi flags trades surrounding several recent events that show win rates and profit concentrations consistent with informed-trading patterns (Mitts & Ofir, 2026). These detection methods are probabilistic rather than ground truth, so they sharpen but do not resolve the count-inflation concern. We do not attempt to construct a de-duplicated account set because the available pseudonymous identifiers on Polymarket (self-chosen display names, sparsely populated bios, wallet creation timestamps) do not partition the wallet population in a way that would bound fragmentation; richer identity-resolution would require on-chain clustering like that developed for cryptocurrency forensics (Meiklejohn et al., 2013). Two complementary analyses elsewhere in the paper partially attenuate this concern. The threshold sensitivity findings (Appendix A.2) are mechanically less affected by singleton fragmentation, since wallets observed in many events are more plausibly individuals than wallets with a single observation. The within-event cross-modal coupling estimates (Figure 6) would understate true person-level coupling if individuals systematically separate trading and commenting across distinct wallets.

Algorithmic trading and market-making bots may similarly inflate trading activity counts. We do not attempt to identify or exclude such accounts because no reliable ground truth for bot-human classification exists in pseudonymous blockchain data. Concentration metrics may be affected: if bots inflate trading counts, the true Gini among human-operated accounts could differ from the values we report. We lack demographic attributes, preventing characterization of who account holders are or how discourse quality varies across accounts. Platform affordances constrain trace availability: we observe only behaviors the API exposes, potentially missing private deliberation, cross-platform coordination, or offline financial decisions that shape observable participation.

Our multi-modal classification assigns accounts to categories based on observed activity in any given modality, thereby representing a conservative upper bound on multi-modal engagement. An account with $50,000 in trades and a single heart reaction is classified the same way as an account with 3,000 comments and $5 in trades. The 4.1% multi-modal figure thus likely overstates the prevalence of multi-modal engagement. We re-estimate the typology under threshold reclassification (requiring at least k events in each modality counted toward the label) and report the resulting partition across $k = \{1, 2, 5, 10, 25, 100\}$ (Appendix A.1). The multi-modal share contracts from 4.1% at $k = 1$ to 1.8% at $k = 10$ and 1.7% at $k = 100$, indicating that the 4.1% figure is itself sensitive to where participation is measured.

Our analyses also aggregate across all 959 election-related events, meaning patterns can be dominated by high-volume markets (primarily the 2024 U.S. presidential election) while masking heterogeneity across smaller markets. Future work should examine event- and market-level disaggregation to determine whether aggregate patterns generalize across individual markets or are driven primarily by the largest ones.

External validity addresses the generalizability of findings to other contexts. This environment reflects a specific user base (crypto-native (Ferguson et al., 2024), predominantly English-language, concentrated on U.S. presidential election markets during 2023–2025) that may not represent other betting platforms (e.g., traditional bookmakers, non-crypto markets), other electoral contexts (e.g., non-U.S. elections, midterms), or broader political communication spaces. The concentration patterns, account type distributions, and cross-mode associations documented here emerged under specific platform design choices (public visibility of positions, no trading fees, prediction market structure rather than parimutuel betting) and regulatory conditions (U.S. CFTC approval in November 2025 (Commodity Futures Trading Commission, 2025)) that differ across jurisdictions and platforms. Empirical assessment is necessary before generalizing these patterns.

Scope constraints further limit interpretation. Our analyses document behavioral patterns that may relate to visibility but do not directly measure visibility itself (e.g., how often accounts view others' trading positions or read comment threads). On Polymarket, position observability is a platform-level constant rather than an account-level variable: all positions are publicly queryable, and all trades are recorded on-chain. Future research could operationalize visibility through platform analytics that track actual viewing behavior, cross-platform comparisons of environments with different default visibility regimes, or experimental designs that manipulate observability in controlled settings. Our focus on account-level participation patterns across behavioral modalities necessarily omits other dimensions of account behavior and discourse. We do not measure trading sophistication (e.g., arbitrage strategies, portfolio diversification, risk-adjusted returns), temporal dynamics within positions (e.g., time-in-position, entry/exit timing relative to events, intraday trading patterns), or the content and quality of political discourse (e.g., argumentation styles, incivility, misinformation, sentiment, linguistic complexity). Natural language processing of comment text could reveal how discourse substance varies by account type, how financial stake correlates with argument quality, or whether multi-modal accounts employ distinct rhetorical strategies. Trading behavior analysis could illuminate whether high-volume traders engage in market-making, whether single-mode traders adopt risk profiles different from those of multi-modal traders, or how position duration relates to commenting activity. We provide foundational account-level descriptions of participation volume and cross-mode engagement; richer characterizations of behavior within modes and content within discourse represent valuable extensions.

Temporal validity poses another constraint (Munger, 2023). The descriptive results presented in the manuscript come entirely from a specific period in the past (May 2023–September 2025), yet the findings will be applied to future settings, where the patterns we document may no longer hold. The rate of change in hybrid political-financial platforms is likely high: platform affordances evolve rapidly (e.g., introduction of new market types, changes to fee structures, algorithmic curation of content), regulatory frameworks shift, and the broader digital ecosystem changes (e.g., emergence of competing platforms, shifts in cryptocurrency adoption, evolution of social media norms around political betting). These dynamics mean that participation patterns documented during the 2024 presidential cycle may not characterize behavior in future elections or on future platforms.

Conclusion

Election betting platforms represent a distinctive class of hybrid digital environments in which financial commitment, political discussion, and social feedback are jointly observable at the wallet level. Through systematic account-level analysis of 778,634 accounts across 846 days, this study documents three foundational patterns: extreme inequality (top 1% generate 36.0–73.7% of activity across modalities; Gini coefficients 0.814–0.951), a systematic modality gradient (financial behaviors most concentrated at 0.951, social reactions intermediate at 0.845, discourse least at 0.814), and within-event decoupling of trading and commenting (both behaviors co-occur in only 21.8% of wallet-event combinations, 90.2% trade-first). For communication scholars, these findings provide a descriptive baseline for studying political discourse in settings where financial positions are visible alongside arguments, with single-mode specialization (95.9%, predominantly Trader Only at 94.4%) and a small multi-modal minority (4.1%) driving disproportionate engagement. For future research on platform governance and regulation, the extreme concentration documented here shows that activity is focused among a narrow segment of highly active accounts, and the modality gradient highlights meaningful differences between financial transactions and discursive acts. As digital platforms increasingly integrate financial incentives into political activities, understanding who participates, how intensely, and across which modes becomes essential for assessing voice and influence. This study provides the empirical foundation for such assessments, establishing baseline patterns against which future platform designs, regulatory frameworks, and comparative contexts can be evaluated, while prioritizing transparent documentation of observable patterns over premature theoretical claims.

Acknowledgements

We are grateful to members of the CODE Lab (Computational Observation of Digital Exposure) and the Center for Communication and Civic Renewal (CCCR) at the University of Wisconsin–Madison for feedback on earlier drafts of this work. We also thank participants of the University of Wisconsin–Madison Google Cloud Platform Research Day for helpful discussion. We thank the editor of the Journal of Quantitative Description: Digital Media for shepherding the manuscript through review.

Funding

This work was supported by the Elections Research Center at the University of Wisconsin–Madison. This work was also funded by the John S. and James L. Knight Foundation and the University of Wisconsin–Madison Office of the Vice Chancellor for Research with funding from the Wisconsin Alumni Research Foundation (Award Reference Number: AWD-002111).

References

Ahn, C., Lelkes, Y., & Levendusky, M. (2024). Sanctioning political speech on social media is driven by partisan norms and identity signaling. *PNAS nexus*, *3*(12), pgae534.

Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *The quarterly journal of economics*, *115*(3), 715–753.

Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? the information content of internet stock message boards. *The Journal of finance*, *59*(3), 1259–1294.

Bartlett, R., & O'Hara, M. (2026). *Adverse selection in prediction markets: Evidence from kalshi* (Working Paper No. 615). Stanford Law School and Cornell SC Johnson College of Business. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6615739

Boulianne, S. (2015). Social media use and participation: A meta-analysis of current research. *Information, communication & society*, *18*(5), 524–538.

Boulianne, S. (2020). Twenty years of digital media effects on civic and political participation. *Communication research*, *47*(7), 947–966.

Brown, A. (2026). Cassandra or the boy who cried wolf? are prediction markets effective early warning systems?

Bürgi, C., Deng, W., & Whelan, K. (2025). *Makers and takers: The economics of the kalshi prediction market* (CESifo Working Paper No. 12122). CESifo Working Papers. https://www.cesifo.org/en/publications/2025/working-paper/makers-and-takers-economics-kalshi-prediction-market

Carr, C. T. (2020). CMC is dead, long live CMC!: Situating computer-mediated communication scholarship beyond the digital age. *Journal of Computer-Mediated Communication*, *25*(1), 9–22.

Chen, H., Duan, X., El Saddik, A., & Cai, W. (2025). Political leanings in Web3 betting: Decoding the interplay of political and profitable motives. *Proceedings of the 17th ACM Web Science Conference 2025*, 96–105.

Christensen, H. S., & Brännlund, A. (2024). Politics on the stock market? political investorism as a form of political participation. *International Political Science Review*, *45*(3), 368–388.

Clinton, J. D., & Huang, T. (2025). *Prediction markets? the accuracy and efficiency of $2.4 billion in the 2024 presidential election* (tech. rep.). Center for Open Science.

Commodity Futures Trading Commission. (2025, November). Amended order of designation: In the matter of the petition of QCX LLC d/b/a Polymarket US to amend its order of designation [Issued November 24, 2025].

Deleep, A., Bai, J., Lee, J., Suresh, D., & Dhawan, H. (2026). How wise is the crowd? bias and edge in prediction markets [arXiv:submit/7314549 [cs.OH]]. *arXiv preprint.*

Della Vedova, J. (2026). *Who profits from prediction markets? execution, not information* (Working Paper No. 6191618). SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6191618

Diercks, A. M., Katz, J. D., & Wright, J. H. (2026, January). *Kalshi and the rise of macro markets* (NBER Working Paper No. 34702). National Bureau of Economic Research. http://www.nber.org/papers/w34702

Douceur, J. R. (2002). The sybil attack. *International workshop on peer-to-peer systems*, 251–260.

Erikson, R. S., & Wlezien, C. (2012). Markets vs. polls as election predictors: An historical assessment. *Electoral Studies*, *31*(3), 532–539.

Ferguson, G., Haglin, K., & Jordan, S. (2024). The personality and politics of cryptocurrency investors. *American Politics Research*, *52*(3), 290–305.

Gerring, J. (2012). Mere description. *British Journal of Political Science*, *42*(4), 721–746.

Gómez-Cram, R., Guo, Y., Jensen, T. I., & Kung, H. (2026). Prediction market accuracy: Crowd wisdom or informed minority? *Available at SSRN 6617059*.

Jain, P., Ezhov, A., Liang, M., Liu, X., Wang, S., Zhou, C., Stoikov, S., & Cetin, U. (2025). Election arbitrage during the 2024 us presidential election. *Available at SSRN 5086104*.

Johansson, W. (2024, June). *Markets vs. machines: An investigation of prediction market forecasting accuracy of economic variables compared to time series forecasting methodologies* [Bachelor's Thesis]. Uppsala University, Department of Economics.

Karpf, D. (2024). Back to basics: Studying digital campaigning while our objects of analysis are in flux. *Political Communication Report*, 1–7.

Liu, S. (2026). *Wisdom of the crowd or wisdom of the insider? insider trading on prediction markets* (Working Paper) (SSRN). University of British Columbia.

Margalit, Y., & Shayo, M. (2021). How markets shape values and political preferences: A field experiment. *American Journal of Political Science*, *65*(2), 473–492.

Marie, A., & Petersen, M. B. (2025). Motivations to connect with like-minded audiences increase partisan sharing on social media. *PNAS nexus*, pgaf197.

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013). A fistful of bitcoins: Characterizing payments among men with no names. *Proceedings of the 2013 conference on Internet measurement conference*, 127–140.

Minozzi, W., Song, H., Lazer, D. M., Neblo, M. A., & Ognyanova, K. (2020). The incidental pundit: Who talks politics with whom, and why? *American Journal of Political Science*, *64*(1), 135–151.

Mitts, J., & Ofir, M. (2026). From iran to taylor swift: Informed trading in prediction markets. *Available at SSRN*.

Munger, K. (2023). Temporal validity as meta-science. *Research & Politics*, *10*(3), 20531680231187271.

Munger, K., Guess, A., & Hargittai, E. (2021). Quantitative description of digital media: A modest proposal to disrupt academic publishing. *Journal of Quantitative Description*, *1*, 1–13.

Muraoka, T., Montgomery, J., Lucas, C., & Tavits, M. (2021). Love and anger in global party politics: Facebook reactions to political party posts in 79 democracies. *Journal of Quantitative Description: Digital Media*, *1*.

O'Brien, E., Coneybeer, J., Boersma, M., & Payne, A. (2023). Political investorism: Conceptualising the political participation of shareholders and investors. *International Political Science Review*, *44*(5), 609–626.

Preda, A. (2013). Tags, transaction types and communication in online anonymous markets. *Socio-Economic Review*, *11*(1), 31–56.

Pyatt, G. (1976). On the interpretation and disaggregation of gini coefficients. *The Economic Journal*, *86*(342), 243–255.

Qiu, L., & Kumar, S. (2017). Understanding voluntary knowledge provision and content contribution through a social-media-based prediction market: A field experiment. *Information Systems Research*, *28*(3), 529–546.

Rabetti, D., Shao, J., & Zhang, C. (2026). *Beating the earnings game: Why do prediction markets outperform professional analysts?* (Working Paper) (First draft: November 2025; this draft: March 2026). SSRN.

Rajadesingan, A., Budak, C., & Resnick, P. (2021). Political discussion is abundant in non-political subreddits (and less toxic). *Proceedings of the International AAAI Conference on Web and Social Media*, *15*, 525–536.

Rhode, P. W., & Strumpf, K. S. (2004). Historical presidential betting markets. *Journal of Economic Perspectives*, *18*(2), 127–142.

Rhode, P. W., & Strumpf, K. S. (2008). Historical political futures markets: An international perspective. *NBER Working Paper*, (w14377).

Sirolly, A., Ma, H., Kanoria, Y., & Sethi, R. (2025). Network-based detection of wash trading. *Available at SSRN*.

Treem, J. W., Leonardi, P. M., & Van den Hooff, B. (2020). Computer-mediated communication in the age of communication visibility. *Journal of Computer-Mediated Communication*, *25*(1), 44–59.

Whelan, K. (2023). On prices and returns in commercial prediction markets. *Quantitative Finance*, *23*(11), 1699–1712.

Whelan, K. (2025). *Agreeing to disagree: The economics of betting exchanges* (tech. rep. No. 126351). Munich Personal RePEc Archive. https://mpra.ub.uni-muenchen.de/126351/

Yang, H.-C. A. (2026). Skilled liquidity provision in prediction markets: Evidence from 150 million trades. *Available at SSRN 6396698*.

Appendix

A.1 *Multi-Modal Reclassification Sensitivity*

The seven-class modality typology in RQ1 classifies a wallet by any observed activity in each modality; this subsection re-estimates the partition under threshold reclassification. At each level of the sensitivity analysis ($k \in \{1, 2, 5, 10, 25, 100\}$), we reclassify each wallet by counting only modalities in which the wallet performs at least k events; wallets falling below k in every modality are dropped from the table. Table 1 reports the resulting per-cohort counts. The multi-modal share contracts from 4.1% at $k = 1$ to 1.8% at $k = 10$ and 1.7% at $k = 100$, while the single-mode share rises correspondingly from 95.9% to 98.3%. Multi-modal cohort sizes attenuate steeply in absolute terms: All Three falls from 10,863 wallets at $k = 1$ to 102 at $k = 100$, and Trader+Commenter from 9,340 to 98. The trajectory thus indicates that the headline 4.1% multi-modal share is sensitive to where the activity threshold is set, and that under any minimum-engagement reclassification, the multi-modal cohorts shrink faster than the single-mode trader cohort.

Table 1: Modality counts under threshold reclassification. At each level k, a wallet's classification counts only modalities in which the wallet performs at least k events. Wallets below k in every modality are excluded. Single-mode share aggregates Trader Only, Commenter Only, and Reactor Only; multi-mode share aggregates the four mixed cohorts.

Modality	$k = 1$	$k = 2$	$k = 5$	$k = 10$	$k = 25$	$k = 100$
Trader Only	734,867	649,972	440,245	293,925	151,619	32,742
Commenter Only	5,470	3,258	1,561	1,119	702	231
Reactor Only	6,121	4,459	2,918	2,136	1,231	493
Trader+Commenter	9,340	4,795	2,040	1,176	468	98
Trader+Reactor	7,169	4,656	2,211	1,213	557	103
Commenter+Reactor	4,804	3,123	1,939	1,331	799	269
All Three	10,863	6,774	3,260	1,739	635	102
Single-mode share	95.9%	97.1%	97.9%	98.2%	98.4%	98.3%
Multi-mode share	4.1%	2.9%	2.1%	1.8%	1.6%	1.7%
Total wallets	778,634	677,037	454,174	302,639	156,011	34,038

A.2 *Threshold Sensitivity of Concentration Statistics*

This subsection reports the threshold sensitivity of the per-account concentration statistics anchored at $k = 1$ in the body. Each level restricts the cohort to accounts with at least k events in the relevant denominator (total activity, trade volume, comments, or reactions). Table 2 reports pooled cohort sizes, Gini coefficients, and top-1% shares across $k \in \{1, 2, 5, 10, 25, 100\}$ for each activity type. Table 3 reports the corresponding per-modality sensitivity analysis for total activity. Table 4 reports the per-modality Gini sensitivity analysis restricted to modalities that include the activity in their classification (e.g., the comment Gini sensitivity analysis is reported for Commenter Only, Trader+Commenter, Commenter+Reactor, and All Three). Across activity types and modalities, Gini values attenuate as k rises but remain substantial; the trade > reaction > comment ordering of pooled Ginis at $k = 1$ persists at every higher level.

Table 2: Pooled threshold sensitivity by activity type. N is the count of accounts at or above k events in the relevant denominator; Gini and Top-1% are computed over that cohort. Trade volume is in U.S. dollars; total activity sums trades, comments, and reactions per account.

Activity	Statistic	$k = 1$	$k = 2$	$k = 5$	$k = 10$	$k = 25$	$k = 100$
Total activity	N	778,634	678,614	457,155	305,350	157,827	34,753
Total activity	Gini	0.850	0.834	0.797	0.769	0.746	0.759
Total activity	Top-1%	0.572	0.564	0.548	0.541	0.541	0.534
Trade volume	N	762,239	666,197	447,756	298,053	153,279	33,045
Trade volume	Gini	0.951	0.945	0.933	0.924	0.916	0.930
Trade volume	Top-1%	0.737	0.729	0.714	0.708	0.704	0.640
Comments	N	30,477	17,950	8,800	5,365	2,604	700
Comments	Gini	0.814	0.765	0.686	0.625	0.540	0.403
Comments	Top-1%	0.359	0.291	0.216	0.174	0.132	0.085
Reactions	N	28,957	19,012	10,328	6,419	3,222	967
Reactions	Gini	0.845	0.805	0.742	0.692	0.620	0.495
Reactions	Top-1%	0.419	0.359	0.286	0.240	0.187	0.113

Table 3: Per-modality threshold sensitivity for total activity. Total activity sums trades, comments, and reactions per account. Modality is held at $k = 1$ classification (recursive convention) so each cohort's composition is fixed across levels.

Modality	Statistic	$k = 1$	$k = 2$	$k = 5$	$k = 10$	$k = 25$	$k = 100$
Trader Only	N	734,867	641,084	428,423	283,536	144,677	29,513
Trader Only	Gini	0.809	0.789	0.741	0.704	0.672	0.696
Commenter Only	N	5,470	2,211	459	165	45	6
Commenter Only	Gini	0.501	0.471	0.463	0.422	0.295	0.085
Reactor Only	N	6,121	3,143	1,239	650	288	109
Reactor Only	Gini	0.885	0.866	0.833	0.796	0.724	0.571
Trader+Commenter	N	9,340	9,340	7,411	5,317	2,665	923
Trader+Commenter	Gini	0.898	0.898	0.883	0.865	0.835	0.767
Trader+Reactor	N	7,169	7,169	5,700	3,977	2,080	728
Trader+Reactor	Gini	0.947	0.947	0.939	0.929	0.908	0.861
Commenter+Reactor	N	4,804	4,804	3,353	2,309	1,283	452
Commenter+Reactor	Gini	0.817	0.817	0.773	0.727	0.656	0.528
All Three	N	10,863	10,863	10,570	9,396	6,789	3,022
All Three	Gini	0.918	0.918	0.916	0.908	0.890	0.850

Table 4: Per-modality Gini sensitivity by activity type. Each row reports the Gini coefficient for a (modality, activity) cell across the threshold sensitivity analysis. Modality cohorts are restricted to those whose classification includes the activity. Modality is held at $k = 1$ classification (recursive convention).

Activity	Modality	$k = 1$	$k = 2$	$k = 5$	$k = 10$	$k = 25$	$k = 100$
Trade volume	Trader Only	0.809	0.789	0.741	0.704	0.672	0.696
Trade volume	Trader+Commenter	0.913	0.906	0.887	0.870	0.837	0.768
Trade volume	Trader+Reactor	0.959	0.955	0.944	0.933	0.913	0.868
Trade volume	All Three	0.953	0.950	0.941	0.930	0.911	0.868
Comments	Commenter Only	0.501	0.471	0.463	0.422	0.295	0.085
Comments	Trader+Commenter	0.514	0.476	0.441	0.421	0.397	0.224
Comments	Commenter+Reactor	0.841	0.803	0.733	0.673	0.585	0.441
Comments	All Three	0.773	0.730	0.658	0.600	0.513	0.375
Reactions	Reactor Only	0.885	0.866	0.833	0.796	0.724	0.571
Reactions	Trader+Reactor	0.809	0.785	0.756	0.728	0.671	0.529
Reactions	Commenter+Reactor	0.846	0.813	0.757	0.705	0.630	0.505
Reactions	All Three	0.796	0.758	0.692	0.636	0.556	0.425

A.3 *Modality by Intensity Contingency*

The Panel B claim that modality and intensity are associated rests on a chi-square test of independence over a 7 × 6 contingency table. We bin accounts on absolute event counts (1, 2–4, 5–9, 10–24, 25–99, 100+) so the bins are mutually exclusive between sensitivity-analysis levels and remain interpretable across modalities. Across all 778,634 accounts, $\chi^2(30) = 45{,}472.59$, $p < .001$, with Cramér's V = 0.108 indicating a small-to-moderate association. The body figure reports the row-percentage view (each cohort's intensity profile, with rows summing to 100%). Here we report the complementary slices: Table 5 reports cell counts, and Table 6 reports column percentages (of accounts in each intensity bin, what fraction belongs to each modality). Trader Only dominates every column, accounting for 93.8% of singleton accounts and 84.9% of accounts with 100+ events; All Three rises from 0.1% of the 2–4 bin to 8.7% of the 100+ bin as multi-modal cohorts gain share at higher activity levels. Multi-modal cohorts have structural floors above $k = 1$ (Trader+Commenter, Trader+Reactor, and Commenter+Reactor require at least two events; All Three requires at least three), producing the zero entries in the singleton column.

Table 5: Modality by intensity contingency: cell counts. Cells report the number of accounts in each (modality, intensity) cell. Intensity bins are absolute event counts per account; modality is the seven-category $k = 1$ classification. Multi-modal cells with structural zeros at $k = 1$ are shown as 0.

Modality	1	2–4	5–9	10–24	25–99	100+	Total
Trader Only	93,783	212,661	144,887	138,859	115,164	29,513	734,867
Commenter Only	3,259	1,752	294	120	39	6	5,470
Reactor Only	2,978	1,904	589	362	179	109	6,121
Trader+Commenter	0	1,929	2,094	2,652	1,742	923	9,340
Trader+Reactor	0	1,469	1,723	1,897	1,352	728	7,169
Commenter+Reactor	0	1,451	1,044	1,026	831	452	4,804
All Three	0	293	1,174	2,607	3,767	3,022	10,863
Total	100,020	221,459	151,805	147,523	123,074	34,753	778,634

Table 6: Modality by intensity contingency: column percentages. Cells report the share of accounts in each intensity bin that belongs to each modality (columns sum to 100% modulo rounding). Multi-modal cohorts have structural zeros in the singleton (1) column.

Modality	1	2–4	5–9	10–24	25–99	100+
Trader Only	93.8%	96.0%	95.4%	94.1%	93.6%	84.9%
Commenter Only	3.3%	0.8%	0.2%	0.1%	0.0%	0.0%
Reactor Only	3.0%	0.9%	0.4%	0.2%	0.1%	0.3%
Trader+Commenter	0.0%	0.9%	1.4%	1.8%	1.4%	2.7%
Trader+Reactor	0.0%	0.7%	1.1%	1.3%	1.1%	2.1%
Commenter+Reactor	0.0%	0.7%	0.7%	0.7%	0.7%	1.3%
All Three	0.0%	0.1%	0.8%	1.8%	3.1%	8.7%

A.4 ***Histograms of Trades, Comments, and Reactions***

The body's CCDF figures use logarithmic axes that compress the lowest-activity bins into a near-vertical wall, so the singleton-and-near-singleton mass is hard to see directly. Figures 7, 8, and 9 report histograms restricted to the low-activity window of 1 to 50 events per account, which covers 91.3% of traders (n = 695,899 of 762,239), 95.4% of commenters (n = 29,087 of 30,477), and 93.8% of reactors (n = 27,166 of 28,957). The trade figure is faceted by modality so the singleton bar's relative height can be read against each cohort's classification floor: single-modality cohorts can have wallets at 1 trade, while multi-modal cohorts have structural minima above 1. The comment and reaction figures are pooled across modalities. In each figure, the leftmost bar is the singleton mass and bars decay rightward, making the long-tailed shape directly visible without the log compression used for the CCDF.

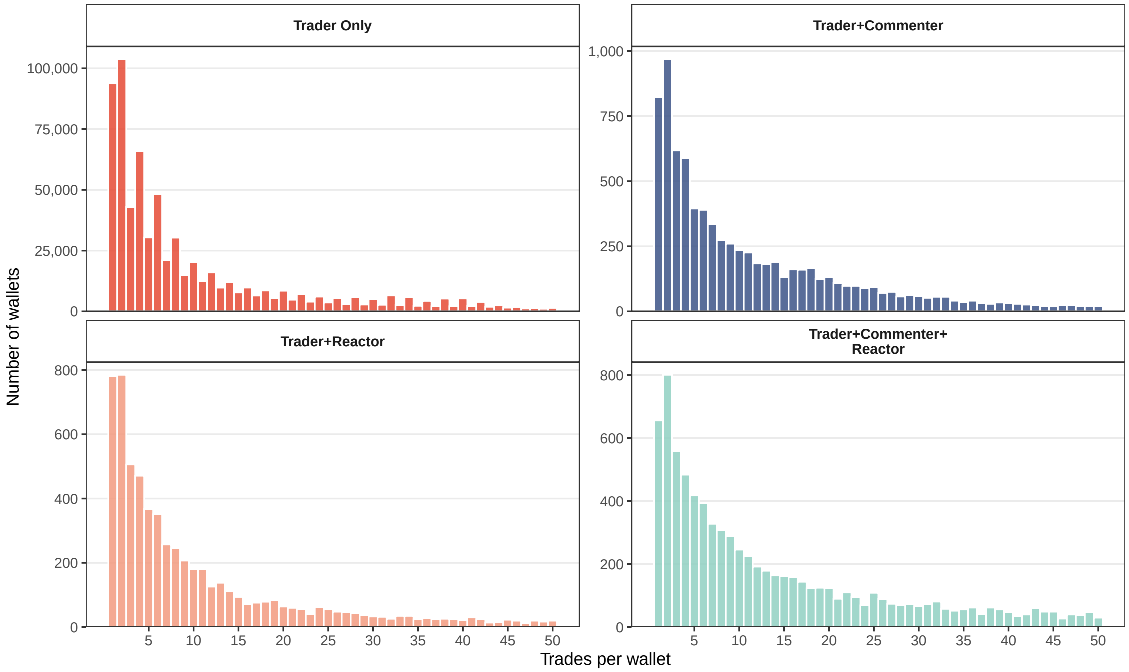

Figure 7. Histogram of trades per account, by modality. Restricted to the 1 to 50 trades window. Single-modality cohorts (Trader Only) include singletons; multi-modal cohorts have structural floors above the lowest bin.

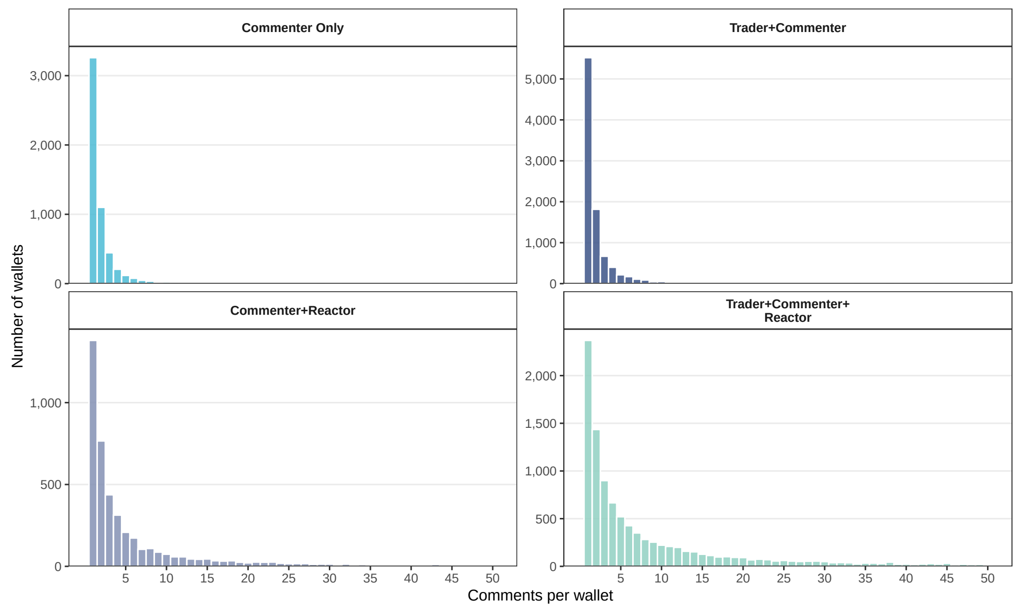

Figure 8. Histogram of comments per commenter. Pooled across modalities; restricted to the 1 to 50 comments window. The singleton bar dominates the distribution.

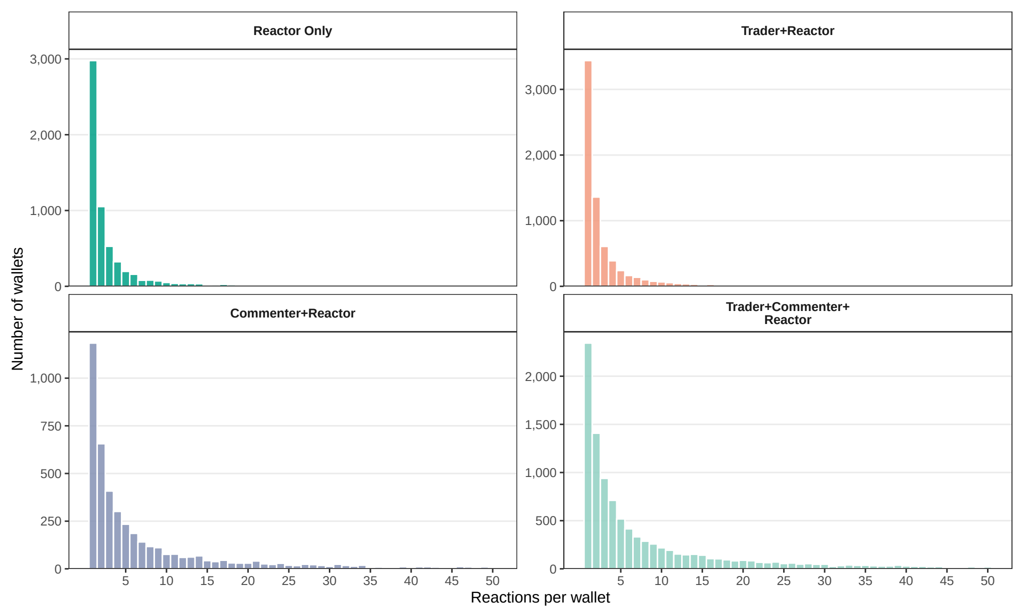

Figure 9. Histogram of reactions per reactor. Pooled across modalities; restricted to the 1 to 50 reactions window. The singleton bar dominates the distribution.

A.5 *Singleton-vs-Active Gini Decomposition*

The body claims that singleton mass and engaged heterogeneity coexist in every modality. To put a number on each contribution, we partition each per-account Gini into a between-group component (the gap between accounts with exactly one event in the modality and accounts with at least two) and a within-group component (heterogeneity among the active subset), following the two-group decomposition of Pyatt (1976). Because the singleton group has zero internal variation by construction, the two components sum exactly to the total Gini. Trade-side decomposition is reported on trade counts so the singleton-vs-active partition is well-defined; the trade-volume Gini cited in the body (0.951) is reported separately. Pooled across all accounts active in a modality (Table 7), between-group inequality accounts for 14.5% of trade-count inequality, 46.5% of comment inequality, and 38.6% of reaction inequality. The per-modality breakdown (Table 8) shows that this share ranges from 6.3% (All Three trade counts) to 71.1% (Commenter Only comments), reflecting how heavily each cohort's inequality is carried by the gap between its singleton accounts and its sustained participants. To assess sensitivity of the decomposition to where the low-vs-high split is drawn, Table 9 reports the same partition across $k_{\text{split}} \in \{2, 5, 10, 25, 100\}$. The between-group share rises mechanically as the cutoff is raised, because more accounts move into the low-activity group and the high-activity group becomes a smaller and more uniform tail; the trajectory documents that singleton-mass and active-heterogeneity contributions are both nonzero across the full range rather than substituting for one another.

Table 7: Pooled singleton-vs-active Gini decomposition by activity type. Each row partitions the per-account Gini into a between-group component (singleton accounts vs. accounts with at least two events in the modality) and a within-group component (heterogeneity among accounts with at least two events). Components sum exactly to the total. Trade decomposition is on trade counts.

Activity	N	Total Gini	Between	Within	% Between	% Within
Trade counts	762,239	0.849	0.123	0.726	14.5%	85.5%
Comments	30,477	0.814	0.378	0.436	46.5%	53.5%
Reactions	28,957	0.845	0.326	0.519	38.6%	61.4%

Table 8: Per-modality singleton-vs-active Gini decomposition. Modality cohorts are restricted to those whose classification includes the activity (e.g., comments are reported for Commenter Only, Trader+Commenter, Commenter+Reactor, and All Three). Modality is held at $k = 1$ classification. The singleton group includes accounts with exactly one event in the activity; the active group includes accounts with at least two.

Activity	Modality	N	Total Gini	Between	Within	% Between	% Within
Trade counts	Trader Only	734,867	0.809	0.123	0.685	15.3%	84.7%
Trade counts	Trader+Commenter	9,340	0.913	0.087	0.826	9.6%	90.4%
Trade counts	Trader+Reactor	7,169	0.959	0.108	0.850	11.3%	88.7%
Trade counts	All Three	10,863	0.953	0.060	0.892	6.3%	93.7%
Comments	Commenter Only	5,470	0.501	0.356	0.145	71.1%	28.9%
Comments	Trader+Commenter	9,340	0.514	0.363	0.150	70.7%	29.3%
Comments	Commenter+Reactor	4,804	0.841	0.275	0.565	32.7%	67.3%
Comments	All Three	10,863	0.773	0.208	0.565	26.9%	73.1%
Reactions	Reactor Only	6,121	0.885	0.455	0.431	51.4%	48.6%
Reactions	Trader+Reactor	7,169	0.809	0.424	0.386	52.3%	47.7%
Reactions	Commenter+Reactor	4,804	0.846	0.238	0.607	28.2%	71.8%
Reactions	All Three	10,863	0.796	0.207	0.589	26.0%	74.0%

Table 9: Singleton-vs-active Gini decomposition across low-vs-high cutoffs. Each row partitions the per-account Gini into a between-group component (accounts with fewer than k_{split} events vs. accounts with at least k_{split}) and a within-group component (heterogeneity among the high-activity group). Components sum exactly to the total Gini. Trade decomposition is on trade counts. The $k_{\text{split}} = 2$ row matches the singleton-vs-active partition in Table 7.

Activity	k_{split}	N	Total Gini	Between	Within	% Between	% Within
Trade counts	2	762,239	0.849	0.123	0.726	14.5%	85.5%
Trade counts	5	762,239	0.849	0.389	0.460	45.8%	54.2%
Trade counts	10	762,239	0.849	0.553	0.296	65.1%	34.9%
Trade counts	25	762,239	0.849	0.670	0.179	78.9%	21.1%
Trade counts	100	762,239	0.849	0.652	0.197	76.8%	23.2%
Comments	2	30,477	0.814	0.378	0.436	46.5%	53.5%
Comments	5	30,477	0.814	0.616	0.198	75.7%	24.3%
Comments	10	30,477	0.814	0.671	0.143	82.4%	17.6%
Comments	25	30,477	0.814	0.653	0.161	80.2%	19.8%
Comments	100	30,477	0.814	0.481	0.333	59.1%	40.9%
Reactions	2	28,957	0.845	0.326	0.519	38.6%	61.4%
Reactions	5	28,957	0.845	0.585	0.260	69.2%	30.8%
Reactions	10	28,957	0.845	0.675	0.170	79.9%	20.1%
Reactions	25	28,957	0.845	0.700	0.145	82.9%	17.1%
Reactions	100	28,957	0.845	0.591	0.255	69.9%	30.1%

A.6 *Joint-Threshold Wallet-Aggregate Trade-Comment Correlation*

The Panel C body claim that the wallet-aggregate trade-comment correlation attenuates but does not collapse under joint thresholding rests on a per-level re-estimation of the Pearson log-log and Spearman rank correlations on the cohort restricted to wallets with at least k trades AND at least k comments. Table 10 reports both correlations with 95% confidence intervals across $k = \{1, 2, 5, 10, 25, 100\}$. The Pearson log-log CI excludes zero at every level from $k = 1$ through $k = 100$, including the smallest cohort (n = 200 at $k = 100$). The Spearman CI excludes zero through $k = 25$ but includes zero at $k = 100$, where the joint-active cohort shrinks below 250 wallets and the rank-based estimator loses precision faster than the parametric estimator on log-transformed counts.

Table 10: Wallet-aggregate trade-comment correlation under joint thresholding. Each row restricts the cohort to wallets with at least k trades AND at least k comments. Pearson log-log r is the Pearson correlation between $\log_{10}(1 + \text{trades})$ and $\log_{10}(1 + \text{comments})$. Spearman r is the rank-correlation companion. 95% confidence intervals use Fisher's z transformation.

k	N wallets	Pearson log-log r	95% CI	Spearman r	95% CI
1	20,203	0.282	[0.269, 0.294]	0.243	[0.229, 0.257]
2	11,569	0.274	[0.257, 0.291]	0.254	[0.236, 0.272]
5	5,300	0.201	[0.175, 0.227]	0.173	[0.145, 0.201]
10	2,915	0.181	[0.146, 0.216]	0.138	[0.100, 0.175]
25	1,103	0.141	[0.082, 0.198]	0.099	[0.037, 0.161]
100	200	0.159	[0.021, 0.292]	0.123	[-0.025, 0.265]

A.7 *Sensitivity of Per-Wallet Joint Trade-and-Comment Realization to Event-Count Cutoff*

The Panel A headline characterizes the per-wallet joint trade-and-comment realization rate among the 20,203 wallets that traded and commented at least once in the observation period. The per-wallet rate is fragile at small numbers of touched events: 5,187 of these wallets (25.7%) touch only a single event, where the per-wallet rate is necessarily 0% or 100%, and a further 7,217 wallets (35.7%) touch only two to four events, where a single event flips the per-wallet rate by 25 to 50 percentage points. To assess whether the headline survives progressively more demanding minimum-event-count thresholds, we report the per-wallet realization rate at four cutoffs (Table 11).

The per-wallet median falls from 40.0% at the inclusive cutoff to 10.0% at the most demanding cutoff (≥ 10 events), reflecting the same long-tailed event-count distribution that produces the histogram's bimodal mass at 0% and 100% in Panel A of the main-text figure. The realization rate pooled across wallet-event combinations (each combination weighted equally regardless of the wallet's total event count) is more stable, ranging from 21.8% at ≥ 1 event to 12.5% at ≥ 10 events. Either summary supports the central claim of Panel A: when a wallet that trades and comments at all in the observation period encounters an event, both behaviors co-occur within that event in only a minority of cases, and that minority shrinks further among the most active subset of the cohort.

Table 11: Per-wallet within-event joint trade-and-comment realization, by minimum-event-count cutoff. Each row restricts the cohort to wallets active in at least k events in the observation period. "Per-wallet median" and "Per-wallet mean" weight each wallet equally; "Pooled rate" weights each wallet-event combination equally and corresponds to total joint events divided by total touched events within the row's cohort.

Cutoff	N wallets	Per-wallet median	Per-wallet mean	Pooled rate
$\geq$ 1 event	20,203	40.0%	48.9%	21.8%
$\geq$ 2 events	15,016	25.0%	31.3%	18.9%
$\geq$ 5 events	7,799	16.7%	19.8%	15.1%
$\geq$ 10 events	3,635	10.0%	14.7%	12.5%

A.8 *Within-Wallet Comment Content in Traded vs. Comment-Only Events*

To assess whether traders comment differently in events they also traded versus events where they only commented, we construct a paired-wallet sample of 2,032 wallets that authored comments in both event strata (1,159 wallets for the per-token Bing sentiment measure, after dropping comments with no sentiment-bearing tokens) and compare per-wallet medians on word count, Flesch reading ease, and Bing sentiment polarity (Table 12).

Table 12: Within-wallet paired comparison of comment content across joint and comment-only events. Each row reports the median per-wallet paired difference (joint events minus comment-only events) on one comment-content measure, computed over the wallets with at least one event in each stratum. "% joint > CO" is the share of paired wallets whose joint-event mean exceeds their comment-only mean. Wilcoxon p is from a paired signed-rank test of the per-wallet paired differences against zero.

Measure	N wallets	Median paired diff.	% joint > CO	Wilcoxon p
Word count	2,032	+0.33	51.2%	0.380
Flesch reading ease	2,032	0.00	50.0%	0.150
Bing sentiment (per token)	1,159	+0.007	51.8%	0.080

A.9 *Nearest-Trade-Delta Clustering of Comments in Joint (Wallet, Event) Pairs*

To assess whether commenting clusters around trading activity within joint (wallet, event) pairs, we compute, for each of the 245,198 comments authored in pairs where the wallet both traded and commented, the signed time difference to the wallet's nearest trade in the same event; quantile and band shares of those deltas are reported in Table 13.

Table 13: Signed time differences from each comment to the same wallet's nearest trade in the same event, across joint (wallet, event) pairs. The sample is 245,198 comments across 31,349 joint pairs (19,252 wallets). Signed deltas are positive when the comment follows the nearest trade and negative when it precedes. The "% before" and "% after" rows partition strictly nonzero signed deltas; ties at zero account for the residual.

Statistic	Value
Median signed delta	+0.20 h
Median absolute delta	4.86 h
Interquartile range, absolute	(0.51, 31.4) h
% within ±1 hour	31.8%
% within ±24 hours	71.1%
% within ±168 hours	92.3%
% before nearest trade	39.4%
% after nearest trade	60.6%

A.10 *Per-Comment Reactions by Same-Event Trading Status*

Whether comments authored by wallets that also traded the same event receive more reactions than comments by wallets that did not is the social-uptake counterpart to the within-event temporal ordering reported in Panel B of Figure 6. The unit of analysis is the comment, not the wallet, and the conditioning is per-(wallet, event) trading rather than the wallet's overall cross-sectional behavior. After restricting to event-attached comments with non-missing wallet, event, and reaction-count fields, the binary contrast partitions 384,576 comments into 245,198 comments authored by wallets that traded the same event (19,252 distinct commenters, 704 distinct events) and 139,378 comments authored by wallets that did not trade that event (13,281 distinct commenters, 627 distinct events).

Comments authored by same-event traders receive systematically more reactions than comments by non-same-event traders (Table 14). The median rises from 0 to 1 reaction, the mean rises from 1.04 to 1.74 reactions, and the share of comments receiving at least one reaction rises from 46.1% to 55.0%. The Mann-Whitney U test on the per-comment reaction count yields a rank-biserial effect size of $r = 0.12$, $p < .001$. The rank-biserial effect size indicates a small-to-modest difference in central tendency: a comment by a same-event trader is more likely than not to receive at least one reaction, while a comment by a non-same-event trader is more likely than not to receive zero. The descriptive contrast does not adjudicate between two consistent interpretations: same-event traders may invest more compositional effort in their comments because they have a financial position riding on the event's outcome, or the audience for an event's comment thread may give more weight to comments by visibly invested traders.

Table 14: Per-comment reactions received, by whether the commenter also traded in the same event. Each row is a per-comment summary across 384,576 event-attached comments retained from the observation period. "% with ≥ 1" is the share of comments in the row's group that received at least one reaction. The Mann-Whitney U test on the underlying per-comment reaction counts yields $r = 0.12$, $p < .001$.

Group	N comments	N commenters	N events	Median	Mean	% with ≥ 1
Same-event trader	245,198	19,252	704	1	1.74	55.0%
Not same-event trader	139,378	13,281	627	0	1.04	46.1%

A.11 *Per-Event Spearman Sensitivity to Minimum-Joint-Wallet Cutoff*

The RQ5 Panel C per-event Spearman correlation between trade volume and comment count is computed at a minimum of 30 joint-active wallets per event. To assess sensitivity of the per-event distribution to this cutoff, we recompute the median per-event Spearman, the median raw-scale Pearson, the median log10-volume Pearson, and the wallet-count-weighted mean Spearman across cutoffs of $k = 10$, 25, 30, 50, and 100 joint wallets per event (Table 15). The central tendency is stable: median Spearman moves between 0.245 and 0.264 across the full range, while the share of joint events that meet the cutoff falls from 40.5% at $k = 10$ to 6.5% at $k = 100$. Raw-scale Pearson is uniformly small (0.07 to 0.10) because per-event whales dominate untransformed volume; log10-volume Pearson runs higher (0.21 to 0.24) and tracks the Spearman ordering. The wallet-count-weighted mean Spearman is similarly flat (0.215 to 0.223), indicating that events contributing more wallets are not driving the median.

Table 15: Per-event correlation summaries across minimum-joint-wallet cut-offs. N is the number of events with at least k joint-active wallets out of 704 total joint events. Spearman is the per-event rank correlation between trade volume and comment count; raw Pearson uses untransformed volume; log Pearson uses $\log_{10}$(volume). The body of the paper reports values at $k = 30$.

k	N	% of joint events	Med. Spearman	Med. raw Pearson	Med. log Pearson	Wgt. mean Spearman
10	285	40.5%	0.259	0.097	0.238	0.223
25	159	22.6%	0.250	0.086	0.227	0.221
30	135	19.2%	0.247	0.080	0.226	0.220
50	94	13.4%	0.245	0.076	0.209	0.217
100	46	6.5%	0.264	0.068	0.210	0.215