Creator Communication and Trading in NFT Markets: Evidence from a Marketplace Policy Change*

Alexey Khazanov† Gleb Kurovskiy‡ Natalia Rostova§

March 2026

Abstract

NFT project creators both inform investors and collect proportional fees on every secondary-market trade. Between 2022 and 2023, platform competition progressively eliminated mandatory creator fees. Using weekly panel data for NFT collections on Ethereum and Solana, we show that creator tweets predict trading volume—primarily through transaction counts rather than prices—and that this link attenuated when fee enforcement weakened, with larger declines for higher-fee collections. We develop a noisy rational expectations model with heterogeneous interpretation of public signals to compare two creator-compensation schemes: per-trade royalties and revenue-equivalent upfront membership fees. Idiosyncratic interpretation noise causes more precise communication to widen belief dispersion rather than compress it, while proportional fees create inaction thresholds that amplify the effect of disagreement on volume. Under royalties, creators back-load disclosure to stimulate trades; under membership fees, they front-load it to resolve uncertainty. The welfare comparison shows that the membership scheme dominates: it eliminates per-trade distortions, and under early-resolution-of-uncertainty preferences the shift in disclosure timing provides an additional welfare gain.

Keywords: NFT markets, information disclosure, creator communication, trading volume, transaction costs

JEL Classification: G12, G14, D82, L86

*We thank the participants of the Hebrew University School of Business Seminar for helpful comments.

†Hebrew University of Jerusalem. Email: alexey.khazanov@mail.huji.ac.il

‡Luminary. Email: gleb.k@luminaryinc.com

§EDHEC Business School. Email: natalia.rostova@edhec.edu

1 Introduction

How does fee structure shape the link between information and trading? We study this question in markets for non-fungible tokens (NFTs)—digital collectibles traded on decentralized platforms. In these markets, project creators are both the primary source of information about their projects and the recipients of proportional fees on every secondary-market transaction. This dual role appears in many financial settings—broker commissions and payment-for-order-flow also tie intermediary income to volume—but NFT markets offer an unusually clean laboratory because all transactions are recorded on public blockchains, fee rates are written into publicly visible code, and competitive platform decisions between 2022 and 2023 produced large, observable changes in whether creator fees were enforced.

Initially, the dominant platform OpenSea automatically deducted contractual creator fees (typically 2.5–10%) from every sale. When new platforms launched offering optional fees, trading shifted to these venues. OpenSea abandoned mandatory enforcement in February 2023, producing a sharp decline in effective creator fees from roughly 5% to near zero while trading technology and asset characteristics remained unchanged.

We combine blockchain transaction records with creator Twitter activity for 252 NFT collections spanning January 2021 through June 2024 across Ethereum and Solana. We construct a weekly panel, estimate two-way fixed effects regressions, decompose volume into transactions versus prices, and use natural language processing to classify tweet content.

One additional creator tweet per week is associated with a 0.79% increase in trading volume on Ethereum and a 1.57% increase on Solana, with both estimates highly significant under two-way fixed effects. The relationship survives first differencing on both blockchains, making it unlikely that common trends in persistent outcomes mechanically drive the association. The response loads on transaction counts rather than per-trade prices—about two-thirds of the volume response operates through the transaction count margin. An hourly event study around individual tweet timestamps shows volume rising within one hour of a tweet, with no pre-tweet ramp-up. Local projections with lag controls confirm that the predictive relationship runs predominantly from tweets to volume; the reverse channel is weaker at every horizon.

The competitive elimination of creator fees occurred in stages. We organize this variation with a three-period design: *pre-erosion* (all platforms enforce mandatory fees), *partial erosion* (new platforms make fees optional while the incumbent maintains enforcement), and *full erosion* (the incumbent abandons enforcement in February 2023). The communication–volume semi-elasticity attenuates when enforcement weakens. Before erosion, collections charging higher fees exhibited steeper communication–volume gradients; after erosion, the fee rate ceases to predict responsiveness. The pattern replicates on Solana despite different market structure. These platform decisions coincided with the November 2022 FTX collapse, so we interpret the estimates as event-time patterns rather than causal effects; the three-period separation, cross-platform replication,

and fee-level heterogeneity collectively support a fee-structure interpretation.

We develop a CARA–Normal noisy rational expectations model augmented with heterogeneous interpretation of public signals to formalize the mechanism, match the empirical patterns qualitatively, and study how welfare improvements from changing fee structure operate primarily by eliminating per-trade distortions, with an additional disclosure-timing channel under early-resolution-of-uncertainty preferences. Investors receive private signals and observe creator communication with idiosyncratic interpretation noise, so that tweets widen cross-sectional belief dispersion rather than compressing it. Proportional fees create inaction thresholds: investors with moderately informative signals do not trade because the fee exceeds their expected gain. Because different investors interpret the same public signal differently, more precise creator communication pushes more investors past the fee-induced inaction threshold, generating three predictions that match the data: larger fees steepen the information–trading relationship, the volume response operates primarily through transaction counts rather than trade size, and creators time disclosure to trigger trades rather than resolve uncertainty. The welfare exercise shows that replacing proportional fees with a revenue-equivalent membership contract improves investor welfare by eliminating the per-trade distortion. In the benchmark upfront membership counterfactual, the fee-structure change also shifts creators' disclosure incentives from late release (to stimulate trading) to early release (to resolve uncertainty); under an early-resolution-of-uncertainty (ERU) preference extension, this behavioral change amplifies the welfare gain. The model nests the standard common-signal REE as interpretation precision grows large, and implies the mechanism operates by aligning information release with efficient uncertainty resolution rather than with trading-volume maximization. This is what makes the erosion of royalty enforcement a useful policy experiment: by changing the trading wedge while leaving creator communication itself in place, platform fee changes reveal how market design affects information transmission, trading activity, disclosure timing, and welfare.

Section 2 reviews related work, Section 3 presents the empirical analysis, Section 4 develops the model, and Section 5 concludes.

2 Related Literature

Our work draws on four literatures: information and trading volume, trading frictions and no-trade regions, social media disclosure, and platform fee design.

2.1 Information, Disagreement, and Trading Volume

We build on the rational-expectations tradition in which heterogeneous information and noise trading generate price informativeness and volume. Grossman and Stiglitz (1980) formalize the role of noise in sustaining trade when prices aggregate information; Admati (1985) provide the canonical noisy REE framework that we extend to include proportional fees; and Kyle (1985)

characterize equilibrium trading with a strategically informed agent. Our theoretical mechanism is closest to disagreement models in which investors are Bayesian but do not condition on the same public object. In the spirit of Harris and Raviv (1993) and Kandel and Pearson (1995), creator communication is publicly observed but heterogeneously interpreted, so more precise disclosure can widen cross-sectional belief dispersion and raise trading activity rather than compress disagreement. This places the paper in the rational-disagreement branch of the literature rather than in behavioral or reduced-form attention models: investors update correctly, but they update on investor-specific interpreted versions of the same creator message.

2.2 Trading Costs and No-Trade Regions

Transaction costs generate inaction regions in which marginal investors do not trade, a classic implication of portfolio choice with frictions (Constantinides, 1986). Frictions and adverse selection shape participation and trade size, and the resulting liquidity effects feed back into volume and prices (Amihud, 2002; Vayanos and Wang, 2012). We focus on *proportional* (percentage) wedges—the dominant fee form in NFT markets—and connect their comparative statics to the observed "information $\times$ cost" interactions.

2.3 Social Media and Disclosure

A growing empirical literature uses textual content and online activity to measure information and attention in asset markets. Tetlock (2007) show that media content predicts returns and trading; Da et al. (2011) use search activity as an attention proxy; Antweiler and Frank (2004) study message boards; Blankespoor et al. (2014) show that firms' Twitter activity affects liquidity. Our setting differs in that the information source is the *creator* rather than dispersed sentiment or formal corporate disclosure, and the same channel can be studied on two blockchains with very different fee environments. On the theory side, the founder's choice of disclosure precision is closest in spirit to Bayesian-persuasion and information-design problems such as Kamenica and Gentzkow (2011), except that here the receiver response is fee-distorted trading volume rather than a static real action.

2.4 NFTs, Blockchain Fees, and Platform Design

NFT markets provide a clean setting: trading is on-chain, royalty rates are contractually specified, and protocol congestion translates into observable transaction costs. Nadini et al. (2021) document key facts about NFT market structure. The economics of blockchain fees is analyzed by Easley et al. (2019) and in work on block-space pricing (Budish, 2025; Huberman et al., 2021). We show that the communication–volume link weakens when royalty enforcement declines and compare per-trade royalties with revenue-equivalent membership fees.

The welfare comparison also connects to two-sided platform pricing (Rochet and Tirole, 2006; Armstrong, 2006; Rysman, 2009). NFT marketplaces sit between creators and traders: per-trade royalties tax marginal transactions and distort information transmission, while fixed-fee alternatives shift extraction away from the trading margin.

2.5 Identification and Inference

We estimate two-way fixed effects specifications and summarize dynamics in event time around the enforcement changes. Because policy timing may correlate with aggregate market conditions, we treat these specifications as event-time diagnostics, not as a clean quasi-experiment. Inference follows Bertrand et al. (2004) and Petersen (2009). Appendix C reports placebo breaks, alternative inference, and robustness for both blockchains.

3 Empirical Analysis

Section 3.2 documents the baseline association between tweets and volume. Section 3.3 examines how this association shifts when royalty enforcement weakens, using a three-regime design that separates Blur's entry from OpenSea's capitulation. Section 3.4 decomposes the response into transactions and prices and examines information content, Section 3.5 replicates on Solana, and Section 3.6 addresses reverse causality.

3.1 Data

We construct a balanced weekly panel of NFT collections spanning January 2021 through June 2024. The data combine transaction records from two blockchain ecosystems (Ethereum and Solana) with creator Twitter activity. All transactions are recorded on public blockchains (immutable distributed ledgers), and contractual fee rates are specified in smart contracts (publicly visible executable code). For each collection we observe weekly trading volume (USD), transaction counts, average prices, and creator-fee rates (typically 2.5–10%). We convert values from native currencies (ETH, SOL) to USD using weekly average exchange rates from Yahoo Finance. Collections are matched to official Twitter accounts, recording weekly tweet counts and content. Appendix A details sample construction.

Table 1 defines the key variables.

The final sample comprises 169 Ethereum collections and 83 Solana collections, each with sufficient trading activity and active Twitter accounts. This yields 17,552 collection-week observations on Ethereum (17,403 in fixed effects regressions after dropping rows with missing identifiers) and 12,450 observations on Solana. Table 2 reports summary statistics.

Table 1: Variable Definitions

Variable	Definition
Volume	Weekly trading volume in USD, summed across all secondary market transactions for the collection
Tweets	Count of tweets posted by the official collection Twitter account in week t
Informativeness	Mean composite informativeness score of the collection's tweets in week t (text-based measure; zero when no tweets)
Announcement Share	Share of weekly tweets classified as announcements (concrete project updates)
Transactions	Number of secondary sale transactions for the collection in week t (Ethereum)
Avg Price	Mean transaction price in USD across all sales in week t (Ethereum); undefined for weeks with zero transactions
Fee Rate	Creator fee rate specified when the collection was created, typically 2.5–10% of transaction value
Regime	Three-level indicator: *Pre-erosion* (week start $<$ 2022-10-25), *Partial erosion* (2022-10-25 $\leq$ week start $<$ 2023-02-21), *Full erosion* (week start $\geq$ 2023-02-21). See Section 3.1.1 for event details.
Post-policy	Indicator equal to 1 for collection-weeks with week start $\geq$ 2022-10-25 (Ethereum) or week start $\geq$ 2022-10-18 (Solana), corresponding to the first full Tuesday–Monday week after the initial royalty-enforcement weakening. Used in two-regime specifications for backward compatibility.

Notes: We convert all USD values from native platform currencies (ETH, SOL) using weekly average exchange rates from Yahoo Finance to avoid mechanical correlation with cryptocurrency price movements. Weeks run Tuesday–Monday.

Table 2: Summary Statistics

Variable	N	Mean	Median	SD	Min	Max
Panel A: Ethereum						
Weekly Volume (USD)	17,552	595,032.98	44,853.53	3,124,173.03	0.07	131,878,739.06
Weekly Transactions	10,005	277.44	72.00	1,136.07	0.00	44,766.00
Average Price (USD)	9,928	4,418.20	878.34	17,908.03	0.02	400,226.65
Weekly Tweet Count	17,552	5.18	0.00	17.08	0.00	451.00
Announcement share	5,158	0.167	0.100	0.210	0.000	1.000
Informativeness score	5,158	0.544	0.534	0.044	0.475	0.765
Royalty rate (%)	169	3.67	3.63	2.49	0.00	9.97
Gas price (gwei)	175	97.3	91.8	29.1	35.6	173.2
Panel B: Solana						
Weekly Volume (USD)	12,450	30,274.30	255.19	98,440.96	0.00	997,838.06
Weekly Tweet Count	12,450	7.92	0.00	24.57	0.00	468.00
Announcement share	3,833	0.151	0.091	0.198	0.000	1.000
Informativeness score	3,833	0.539	0.529	0.042	0.485	0.860
Royalty rate (%)	83	1.12	1.10	0.37	0.72	4.06

Notes: Announcement share and informativeness are computed on tweet-active weeks (tweet_count $>$ 0). Transactions and Average Price are reported only for weeks with positive trading volume. Creator fee rates are collection-level (one observation per collection). Gas prices on Ethereum (network congestion costs) are weekly averages; Solana's gas costs are negligible by comparison.

The gap between median ($44,854) and mean ($595,033) weekly Ethereum volume reflects the heavy right tail characteristic of NFT trading. The median collection posts zero tweets in a given week, with occasional bursts. Both distributions favor the log-linear specifications used throughout. Appendix Figures 17–20 report distributions and time series for both blockchains.

3.1.1 Royalty Enforcement and Marketplace Competition

Creator fees in NFT markets are not automatically enforced by the underlying platform technology. Instead, enforcement depends on the policies of individual trading venues (marketplaces). When a buyer purchases an NFT on a marketplace that enforces creator fees, the marketplace automatically deducts the specified percentage from the transaction and transfers it to the project creator; when the same asset trades on a marketplace that does not enforce these fees, the creator receives nothing. This creates competition among marketplaces: those that enforce fees provide revenue to creators but may lose trading volume to zero-fee competitors.[1]

Between July 2022 and February 2023, competitive pressure among trading venues progressively eliminated creator-fee enforcement on Ethereum. The sequence unfolded as follows:

1. **SudoSwap launch (July 2022).** SudoSwap, an automated market maker for NFTs,

[1]Solana introduced technology-enforced creator fees for new collections in late 2022, but adoption was limited and the dominant marketplace (Magic Eden) still controlled enforcement through its own policies.

launched with zero creator royalties by design, reaching $50 million in cumulative volume within two months.[2]

2. **X2Y2 optional royalties (August 2022).** X2Y2 allowed buyers to choose whether to pay creator royalties at 100%, 50%, or 0% of the contractual rate.[3]

3. **Magic Eden shift on Solana (October 15, 2022).** Magic Eden, holding 90% market share on Solana, moved to buyer-optional creator fees and simultaneously dropped its own 2% platform fee.[4]

4. **Blur launch (October 19, 2022).** Blur entered the Ethereum market with zero platform fees, optional creator fees, and a token airdrop program designed to attract trading volume.[5]

5. **LooksRare shift (October 27, 2022).** LooksRare made creator-fee payments opt-in, replacing them with a mandatory 0.5% platform fee that would be distributed to creators.[6]

6. **OpenSea enforcement tool (November 6–8, 2022).** OpenSea deployed an on-chain tool allowing new collections to block trades on marketplaces that did not honor creator fees, though adoption was opt-in.[7]

7. **OpenSea capitulation (February 17, 2023).** OpenSea (i) dropped its 2.5% platform fee to zero, (ii) made creator fees optional with a 0.5% minimum for collections not using the enforcement tool, and (iii) stopped blocking competing marketplaces.[8] This was the largest single shift in effective enforcement given OpenSea's 80% market share.[9]

8. **Enforcement tool sunset (August 2023).** OpenSea announced it would permanently disable its fee-enforcement tool, effective August 31, 2023.[10]

We organize this sequence into three enforcement regimes based on when effective creator fees declined. The *pre-erosion* period (before October 25, 2022) covers the era of full enforcement, when the dominant marketplace (OpenSea, holding 80% of trading volume) automatically deducted and remitted creator fees. The *partial erosion* period (October 25, 2022 through February 20, 2023) begins when the major new entrant launched with optional fees, creating a competitive

[2] See "The Case for 0% Royalty Fees: Sudoswap vs Opensea," *Medium*, 2022; "Royalty-free NFT platform SudoAMM surpasses $50 million in volume," *W3B Portfolio News*, September 2022.

[3] "X2Y2 implements flexible royalty option for buyers," *CryptoSlate*, August 26, 2022.

[4] "NFT Marketplace Magic Eden Moves to Optional Royalty Model," *CoinDesk*, October 15, 2022.

[5] "An Introduction to the New NFT Marketplace Blur," *NFT News Today*, October 27, 2022.

[6] *The Defiant*, October 2022; LooksRare documentation, 2022.

[7] "OpenSea launches on-chain tool to enforce NFT royalties," *Cointelegraph*, November 2022.

[8] "OpenSea Just Cut Fees and Creator Royalties," *NFT Now*, February 17, 2023; "OpenSea Goes Zero-Fee, Creator Royalties Optional," *Yahoo Finance*, February 2023.

[9] "Upstart NFT marketplace Blur has left OpenSea in the dust—for now," *Fortune*, February 26, 2023.

[10] "NFT marketplace OpenSea to disable royalty enforcement tool," *Cointelegraph*, August 2023.

environment in which some venues enforced fees while others did not. The *full erosion* period (February 21, 2023 onward) begins when the incumbent abandoned mandatory enforcement, effectively ending creator-fee collection across all major venues on the Ethereum platform.

Transaction data confirm the timeline: the volume-weighted effective creator-fee rate in our sample (calculated as total fees collected divided by total trading volume) stood at 4.4% in the week of the new entrant's launch, fell to 2.7% within two weeks as trading volume shifted to zero-fee venues, and collapsed to 0.6% after the incumbent capitulated. Figure 1 visualizes this enforcement erosion over time. On the Solana platform, we use a single policy break corresponding to the dominant marketplace's October 2022 shift to optional fees, which effectively determined the enforcement regime given that marketplace's 90% share of Solana trading volume.

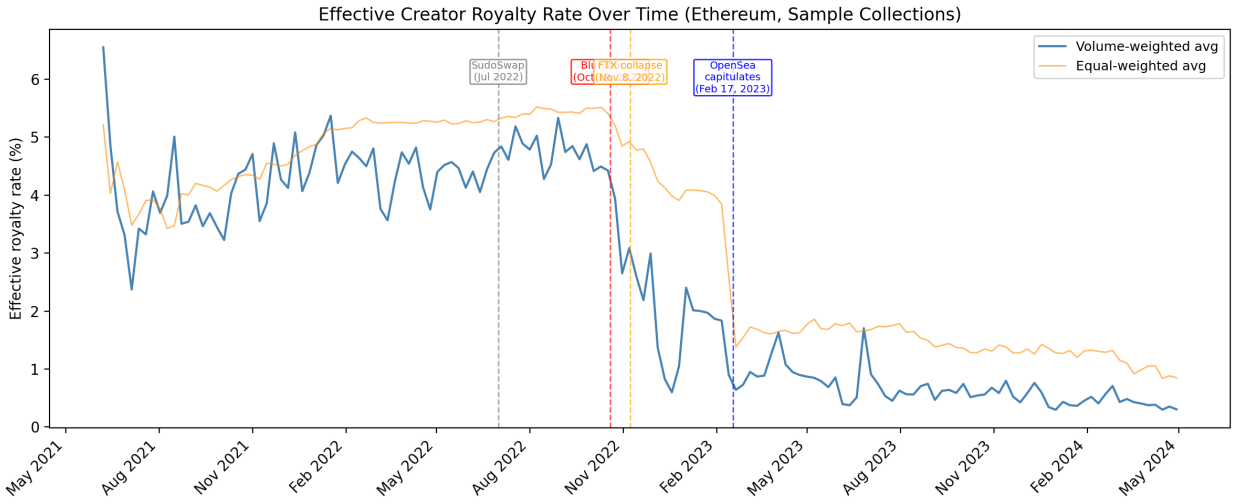

Figure 1: Effective Creator Royalty Rate Over Time (Ethereum)

Notes: Volume-weighted effective royalty rate (creator fees collected / trading volume) computed weekly across sample collections. Vertical lines mark key enforcement events: SudoSwap launch (July 2022), Blur launch with optional fees (October 19, 2022), FTX collapse (November 8, 2022), and OpenSea capitulation to optional fees (February 17, 2023). The rate drops from over 4% before Blur's entry to below 1% after OpenSea abandoned enforcement.

3.2 Baseline Results

The baseline specification is:

$$\log(1 + \text{Volume}_{it}) = \alpha_i + \delta_t + \beta \, \text{Tweets}_{it} + \varepsilon_{it}, \tag{1}$$

where α_i and δ_t are collection and week fixed effects, and Tweets_{it} is the raw tweet count (not log-transformed). The coefficient β is a semi-elasticity: it measures the percentage increase in volume associated with one additional tweet. Standard errors are two-way clustered by collection

and week throughout.

Tables 3 and 4 build up in stages: pooled OLS (columns 1–2), two-way fixed effects (column 3), lagged volume (columns 4–5), and first differences (columns 6–7).

Table 3: Creator Communication and Trading Volume (Ethereum)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tweet count	0.0129***	0.0123***	0.0079***	0.0033***	0.0033***	0.0045***	0.0045***
	(0.0034)	(0.0029)	(0.0023)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
Collection FE	No	No	Yes	Yes	Yes	No	No
Week FE	No	No	Yes	Yes	Yes	Yes	Yes
Lagged volume	No	No	No	Yes	Yes	No	No
Market controls	No	Yes	No	No	Yes	No	Yes
Observations	17,403	17,403	17,403	17,234	17,234	17,234	17,234

Notes: Dependent variable is log(1+Volume). The regressor is tweet count (not log-transformed), so coefficients are semi-elasticities. Standard errors are two-way clustered by collection and week. Appendix Table 29 reports the log–log elasticity specification as a robustness check. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Creator Communication and Trading Volume (Solana)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tweet count	0.0321***	0.0265***	0.0157***	0.0042***	0.0042***	0.0048**	0.0048**
	(0.0065)	(0.0067)	(0.0055)	(0.0012)	(0.0012)	(0.0023)	(0.0023)
Collection FE	No	No	Yes	Yes	Yes	No	No
Week FE	No	No	Yes	Yes	Yes	Yes	Yes
Lagged volume	No	No	No	Yes	Yes	No	No
Market controls	No	Yes	No	No	Yes	No	Yes
Observations	12,450	12,450	12,450	12,367	12,367	12,367	12,367

Notes: Same specification as Table 3 applied to the Solana sample. Standard errors are two-way clustered by collection and week. Appendix Table 30 reports the log–log elasticity specification. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Under two-way fixed effects (column 3), the semi-elasticity is 0.0079 on Ethereum and 0.0157 on Solana, both highly significant. This means that one additional tweet per week is associated with a 0.79% increase in trading volume on Ethereum and a 1.57% increase on Solana. Adding lagged volume attenuates the coefficient (column 4), indicating that part of the comovement reflects persistence. First differencing (column 6) yields comparable short-run semi-elasticities (0.0045 on both blockchains); the levels relationship is not an artifact of common trends. At the median weekly volume of $45,000 on Ethereum, the baseline semi-elasticity of 0.0079 corresponds to roughly $355 per tweet. Across the 169 Ethereum collections over the sample period, the implied aggregate association is economically meaningful, though this reflects correlation rather than a causal effect.

3.3 Policy Events

The baseline estimates are silent about *why* tweets and volume co-move. A series of competitive marketplace decisions between July 2022 and February 2023 (Section 3.1.1) progressively weakened royalty enforcement on Ethereum. Two events bracket the sharpest decline: Blur's launch on October 19, 2022, which introduced a zero-fee, optional-royalty competitor, and OpenSea's capitulation on February 17, 2023, when the dominant marketplace (∼80% market share) dropped mandatory enforcement and cut its platform fee to zero. On Solana, Magic Eden moved to optional royalties on October 15, 2022. None of these events is plausibly exogenous to aggregate market conditions—the FTX collapse occurred between the two Ethereum events—so the estimates below are best read as event-time diagnostics rather than causal effects. The three-regime design, together with cross-chain replication and fee-rate heterogeneity, disciplines the wedge interpretation.

We begin with a pre/post split at Blur's launch:

$$\begin{aligned}\log(1+\text{Volume}_{it}) = \alpha_i + \delta_t \\ + \beta_{\text{Pre}}\,\mathbf{1}\{t < \text{Blur}\}\,\text{Tweets}_{it} \\ + \beta_{\text{Post}}\,\mathbf{1}\{t \geq \text{Blur}\}\,\text{Tweets}_{it} + \varepsilon_{it}.\end{aligned} \tag{2}$$

To separate Blur's entry from OpenSea's capitulation, we also estimate a three-regime specification (Section 3.1.1):

$$\begin{aligned}\log(1+\text{Volume}_{it}) = \alpha_i + \delta_t \\ + \beta_{\text{Pre}}\,\mathbf{1}\{t \in \text{Pre}\}\,\text{Tweets}_{it} \\ + \beta_{\text{Part}}\,\mathbf{1}\{t \in \text{Partial}\}\,\text{Tweets}_{it} \\ + \beta_{\text{Full}}\,\mathbf{1}\{t \in \text{Full}\}\,\text{Tweets}_{it} + \varepsilon_{it},\end{aligned} \tag{3}$$

where *Pre-erosion* denotes weeks before October 25, 2022; *Partial erosion* covers October 25, 2022 through February 20, 2023 (Blur and competitors offer optional royalties; OpenSea maintains enforcement); and *Full erosion* denotes February 21, 2023 onward (OpenSea drops mandatory enforcement).

Event-time plots (Figure 2; Appendix Figures 21 and 22) reveal substantial time variation. On Ethereum the semi-elasticity is roughly flat before Blur and drops at the policy week. On Solana the pre-period is noisier, but the trough again coincides with the policy date. Both panels show a temporary rebound around the FTX collapse. Appendix Figures 31 and 32 replicate the event-study design using the raw tweet count (semi-elasticity specification) rather than the log-log specification, confirming that the attenuation pattern is not an artifact of the functional form.

Table 5 reports the two-regime semi-elasticities from equation (2).

The semi-elasticity is 0.0103 in the pre-erosion period, meaning one additional tweet increased

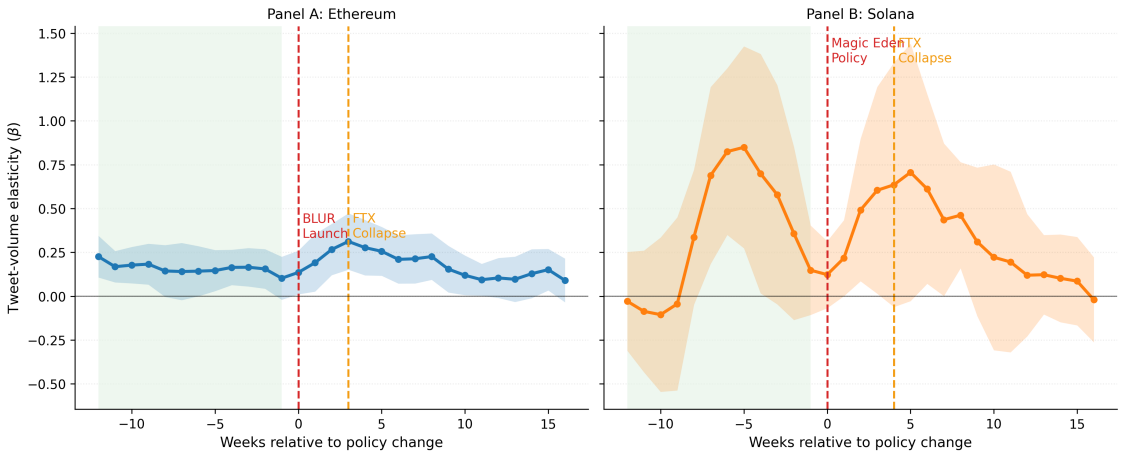

Figure 2: Cross-Blockchain Comparison: Event Studies Around Policy Changes

Notes: Rolling 8-week window TWFE estimates of the semi-elasticity coefficient on Tweets_{it} in equation (1). Each point uses weeks $k-3$ through $k+4$ relative to the policy week, with collection and week fixed effects; shaded areas are 95% confidence intervals based on two-way clustering by collection and week. Week 0 corresponds to Blur's launch (Ethereum, October 19, 2022) and Magic Eden's shift to optional royalties (Solana, October 15, 2022). The vertical red dashed line marks week 0. The vertical orange dashed line marks the FTX collapse (November 8, 2022). The vertical green dashed line (Ethereum panel only) marks OpenSea's capitulation (February 17, 2023), when the dominant marketplace dropped mandatory enforcement. The light green band highlights the pre-policy window (-12 to -1).

Table 5: The Tweet-Volume Relationship Before and After Enforcement Weakening (Ethereum)

	(1)	(2)
Tweet count	0.0103**	0.0043***
	(0.0045)	(0.0014)
Tweet count × Post-policy	-0.0047	-0.0018
	(0.0050)	(0.0015)
Collection FE	Yes	Yes
Week FE	Yes	Yes
Lagged volume	No	Yes
Observations	17,403	17,234

Notes: Dependent variable is log(1+Volume). The regressor is tweet count (not log-transformed) interacted with a post-policy dummy. The main coefficient is the pre-erosion semi-elasticity; the interaction captures the change post-erosion. Post-erosion semi-elasticity = main coefficient + interaction. Standard errors are two-way clustered by collection and week. Appendix Table 31 reports the log–log elasticity specification. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

volume by 1.03%. After enforcement weakened, the semi-elasticity drops to 0.0056 (= 0.0103 - 0.0047), implying one tweet increases volume by only 0.56%—a decline of 0.47 percentage points, or roughly a 46% reduction in the marginal effect. While the point estimates suggest substantial attenuation, the change is not statistically significant at conventional levels ($p = 0.35$). The event-time pattern (Figure 2) and the Solana evidence (Section 3.5) reinforce this reading.

3.3.1 Heterogeneity by Royalty Rates

If the wedge matters, the attenuation should concentrate among collections with higher contractual royalty rates. We interact tweet count with the royalty rate and the weak-enforcement indicator.

In Table 6, the semi-elasticity is increasing in the royalty rate before enforcement weakening, and the weak-enforcement interaction is negative, focusing the decline on high-royalty collections. Figure 3 plots the implied semi-elasticities across the royalty distribution.

Table 6: Fee Heterogeneity Around Enforcement Weakening (Ethereum)

	(1)
Tweet count	0.0104**
	(0.0045)
Tweet count × Fee rate	0.3617
	(0.2612)
Tweet count × Weak enforcement	-0.0054
	(0.0046)
Tweet count × Weak enforcement × Fee rate	-0.4352*
	(0.2545)
Collection FE	Yes
Week FE	Yes
Observations	17,403

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with royalty rate and weak-enforcement indicator. Coefficients represent percentage change in volume per tweet at mean royalty rate. Interaction coefficients represent the change in semi-elasticity per percentage point increase in royalty rate. Standard errors are two-way clustered by collection and week. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Table 7 extends the fee-heterogeneity analysis to three regimes. The fee interaction is positive in the pre-erosion period (0.36, though imprecisely estimated), turns negative during partial erosion (-0.19), and is indistinguishable from zero under full erosion (-0.02)—a pattern of diminishing fee sensitivity as the wedge loses bite, though none of the individual interaction coefficients reach conventional significance levels.

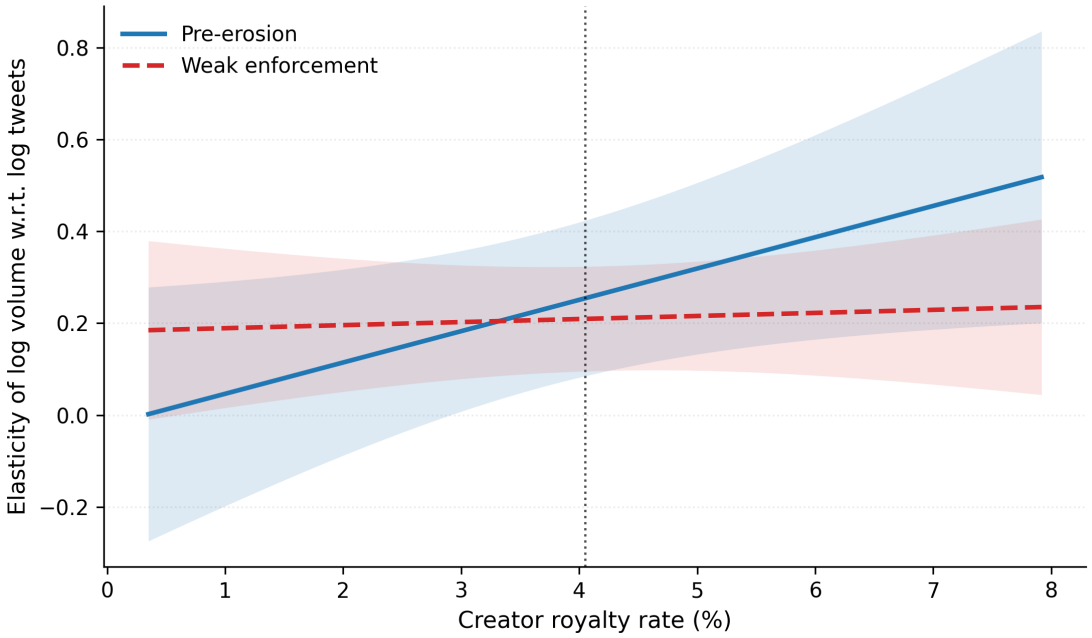

Figure 3: Tweet-Volume Semi-Elasticity by Royalty Rate (Ethereum)

Notes: Lines plot the implied semi-elasticity (percentage change in volume per additional tweet) from Table 6 as a function of the contractual royalty rate. Shaded bands show 95% confidence intervals computed using the delta method.

Table 7: Fee Heterogeneity Across Enforcement Regimes (Ethereum)

	Pre-erosion	Partial erosion	Full erosion
Tweet count	0.0104**	0.0049*	0.0058***
	(0.0045)	(0.0026)	(0.0021)
Tweet count × Fee rate	0.3619	-0.1915	-0.0218
	(0.2635)	(0.1521)	(0.0841)
p-value: Pre vs. Full (interaction)		0.138	
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations		17,403	

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with the demeaned contractual royalty rate (Full_CF_rate minus its sample mean). Coefficients represent percentage change in volume per tweet. Interaction coefficients represent the change in semi-elasticity per percentage point increase in royalty rate. Standard errors are two-way clustered by collection and week.

3.3.2 Network Congestion (Gas Fees)

Royalties are not the only transaction cost on Ethereum; network congestion raises gas fees, which vary week to week and are absorbed by the week fixed effects. We interact gas with collection-level communication to ask whether higher costs steepen the within-week information–volume gradient. The test targets the slope, not the level: does communication matter more when trading is more expensive?

Table 8 interacts tweet count with (demeaned) log gas price and the weak-enforcement indicator. The coefficient on the interaction is positive: the semi-elasticity rises in high-gas weeks. Figure 4 plots the implied semi-elasticities across the gas distribution. The Online Appendix maps royalties and gas into the model's per-trade wedge $\Delta_{i,t}$: on Ethereum the median gas component is comparable to the median royalty wedge; on Solana it is negligible.

Table 8: Congestion Costs and the Tweet-Volume Semi-Elasticity (Ethereum)

	(1)
Tweet count	0.0090**
	(0.0040)
Tweet count × Weak enforcement	-0.0032
	(0.0047)
Tweet count × (Log Gas − Mean)	0.0129*
	(0.0077)
Tweet count × Weak enforcement × (Log Gas − Mean)	-0.0080
	(0.0083)
Collection FE	Yes
Week FE	Yes
Observations	17,403

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with log gas price (demeaned) and weak-enforcement indicator. Coefficients represent percentage change in volume per tweet at mean gas price. Interaction coefficients represent the change in semi-elasticity per unit change in log gas price. Standard errors are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Separating the gas analysis by enforcement period (Table 9), gas amplification of the semi-elasticity persists in both pre-erosion and full-erosion periods but reverses sign during partial erosion. The sign flip during partial erosion likely reflects transitional trading dynamics around the Blur/FTX period, when both fee structures and aggregate market conditions shifted simultaneously. Figure 5 plots the implied semi-elasticity curves for each regime.

Gas amplifies *information*-relevant communication more than generic activity. The Online Appendix decomposes tweets into announcements and other communication and interacts each with gas: the announcement–volume slope rises sharply with congestion while the slope on other tweets

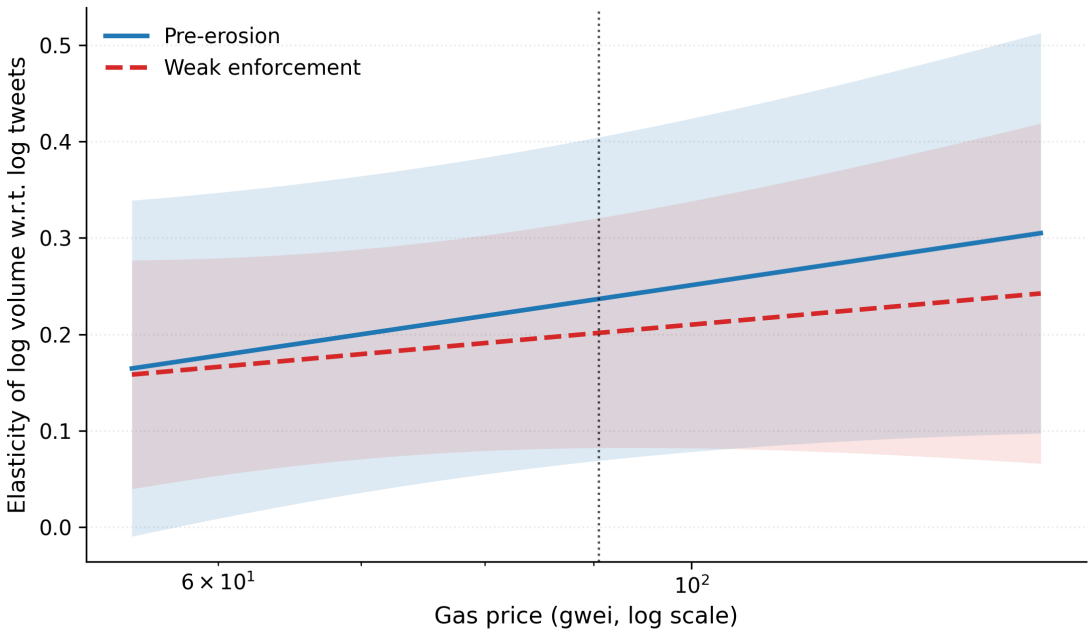

Figure 4: Tweet-Volume Semi-Elasticity by Network Congestion (Ethereum)

Notes: Lines plot the implied semi-elasticity from Table 8 as a function of the Ethereum gas price (in gwei). Shaded bands show 95% confidence intervals computed using the delta method.

Table 9: Congestion Costs and the Tweet-Volume Semi-Elasticity Across Regimes (Ethereum)

	Pre-erosion	Partial erosion	Full erosion
Tweet count	0.0091**	0.0014	0.0064***
	(0.0040)	(0.0028)	(0.0020)
Tweet count × (Log Gas − Mean)	0.0128*	-0.0132***	0.0093**
	(0.0077)	(0.0048)	(0.0047)
p-value: Pre vs. Partial (gas interaction)		0.001	
p-value: Pre vs. Full (gas interaction)		0.685	
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations		17,403	

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with demeaned log gas price in each regime. Coefficients represent percentage change in volume per tweet. Wald tests compare the gas interaction coefficient across regimes. Standard errors are two-way clustered by collection and week.

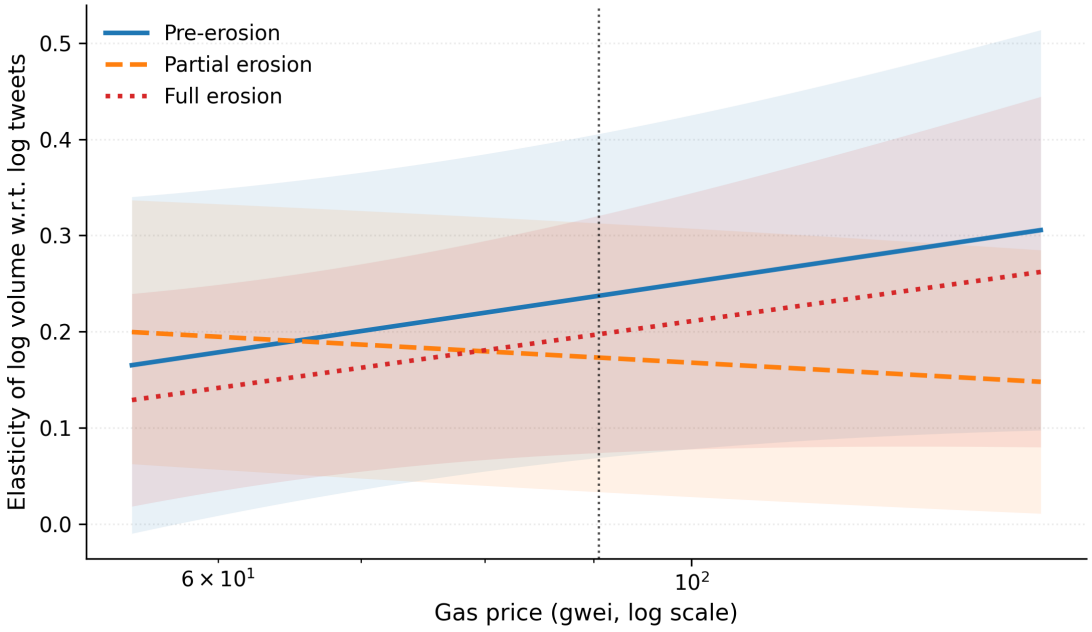

Figure 5: Tweet-Volume Semi-Elasticity by Gas Price Across Enforcement Regimes (Ethereum)

Notes: Lines plot the implied semi-elasticity from Table 9 as a function of gas price (gwei) for each enforcement regime. Shaded bands show 95% confidence intervals computed using the delta method.

is flat. A text-based informativeness score yields the same pattern, and the amplification fades under weak enforcement. Because gas varies within weeks, these interactions are not confounded by the Blur/FTX timing.

3.3.3 The FTX Confound

FTX collapsed on November 8, 2022—three weeks after Blur—and the two shocks cannot be cleanly separated. Several patterns help narrow the interpretation: the event-time drop begins at Blur's launch, before FTX (Appendix Figure 21); the same attenuation appears on Solana around Magic Eden's contemporaneous policy change, on a blockchain with different architecture and market composition; and our outcome is a within-collection *slope*, not a level—FTX depressed volume everywhere, but a common level shock does not mechanically flatten the gradient. No single piece of evidence is decisive, but together they make a pure aggregate-shock explanation less compelling.

The three-regime specification adds further separation. The FTX collapse occurred in November 2022, squarely within the partial-erosion window. In Table 10 (equation 3) the semi-elasticity drops from 1.03% (pre-erosion) to 0.44% (partial) and 0.61% (full erosion). The largest decline coincides with Blur's entry, and the coefficient stays down even after markets partially recovered from FTX—the full-erosion estimate is statistically indistinguishable from the partial-erosion estimate ($p = 0.596$), but both sit below the pre-erosion level. That the attenuation persists three

months after FTX, precisely when OpenSea abandoned enforcement, is difficult to attribute to an aggregate shock alone.

Two additional confounds deserve mention. First, Blur's multi-season airdrop program (Season 1: October 2022–February 2023; Season 2: February 2023 onward) rewarded trading activity through loyalty multipliers and bid-to-earn mechanics.[11] The resulting inorganic activity was substantial: weekly volume in our sample jumped from roughly $80–115 million in the weeks before the February 14 token launch to $361 million in the airdrop week and $808 million the following week—a 3.8-fold spike that faded over the next month. Third-party estimates of wash trading on Blur range from roughly 14% of cumulative volume, using wallet-matching heuristics, to over 80% of recorded volume in the week following the airdrop, using a broader "inorganic trading" criterion.[12] These trades are unlikely to respond to creator communication, potentially attenuating the tweet–volume semi-elasticity even if the underlying information channel is unchanged. Second, Blur attracted a more concentrated, sophisticated trader base—roughly 80% of its volume came from approximately 500 wallets[13]—and the volume share of high-royalty collections (above the sample median of 3.6%) declined steadily from 51% in 2022Q1 to 21% in 2023Q4. This compositional shift toward low-fee collections could independently weaken the average semi-elasticity.

In principle, both channels could attenuate the semi-elasticity even if the underlying information channel were unchanged. However, neither channel predicts the *fee-heterogeneity* pattern: wash trading and compositional shifts should reduce the semi-elasticity uniformly across fee levels (absent systematic correlation with royalty rates), yet the triple interaction in Table 7 shows that the attenuation is concentrated among high-royalty collections. Similarly, the announcement-specific decline (Tables 14 and 16) is harder to attribute to inorganic volume, which should attenuate information-rich and information-poor tweets alike. We also note that FTX and Alameda Research were deeply embedded in the Solana ecosystem, holding approximately 58 million SOL tokens,[14] so the Solana replication in Section 3.5 does not provide a fully independent cross-chain test. Table 37 in the Appendix shows that the baseline results are robust to excluding the Blur airdrop window (February 14–March 14, 2023), restricting to the pre-erosion subsample only, and winsorizing volume at the 99th percentile. Table 35 reports a specification that replaces the discrete regime dummies with a continuous effective-fee proxy that interacts the contractual creator-fee rate with regime-specific enforcement factors.

[11] Galaxy Research characterized the resulting activity as "inorganic"; see "Blur's Airdrop Shaped Inorganic Surge in 2023 NFT Trading," *Blockworks*, 2023.

[12] Hildobby, "NFT Wash Trading," Dune Analytics dashboard, February 2023 (14%, flagging same-wallet trades and repeated purchases). CryptoSlam flagged $577 million in inorganic trades; see "Wash Trading on Blur Has NFT Community Up In Arms," *Decrypt*, February 21, 2023, and "CryptoSlam Flags Blur Wash Trades," *Forkast*, February 2023.

[13] DappRadar data as reported in "Upstart NFT marketplace Blur has left OpenSea in the dust—for now," *Fortune*, February 26, 2023.

[14] Solana Foundation, "Solana Foundation and Solana Labs Facts Related to the FTX Bankruptcy," November 2022 (58.06 million SOL across five transactions).

Table 10: The Tweet-Volume Relationship Across Enforcement Regimes (Ethereum)

	Pre-erosion	Partial erosion	Full erosion
Tweet count	0.0103**	0.0044	0.0061***
	(0.0045)	(0.0031)	(0.0020)
p-value: Pre vs. Partial		0.282	
p-value: Pre vs. Full		0.413	
p-value: Partial vs. Full		0.596	
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations		17,403	

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with three mutually exclusive regime indicators: pre-erosion (before October 25, 2022), partial erosion (October 25, 2022 to February 20, 2023), and full erosion (from February 21, 2023). Coefficients represent percentage change in volume per additional tweet. Wald tests report p-values for pairwise equality of regime coefficients. Standard errors are two-way clustered by collection and week.

3.3.4 Interpretation

We interpret the four patterns—semi-elasticity decline, cross-chain replication, fee heterogeneity, and congestion amplification—as consistent with a single mechanism. While the timing of marketplace policy changes coincided with broader market events (notably the FTX collapse), we do not claim quasi-experimental identification. Instead, the convergence of cross-platform replication, differential attenuation by fee level, interaction with network congestion, and the concentration of effects on the extensive margin collectively point toward a fee-wedge mechanism rather than a pure aggregate shock.

A mechanism consistent with the empirical patterns operates as follows. Proportional fees create an inaction region: only investors with sufficiently strong signals trade, because the wedge must be covered by expected gains. When investors interpret the same creator communication with idiosyncratic noise—as in the heterogeneous-interpretation model of Section 4—more precise public signals widen the cross-sectional dispersion of posterior beliefs rather than compressing it. This wider dispersion pushes more investors past the fee-induced inaction threshold, which is why the tweet-volume semi-elasticity is positive in the first place. Higher royalties widen the inaction band and steepen the gradient (Table 7); weaker enforcement shrinks the wedge and flattens it (Table 10); network congestion acts as an additional wedge component (Table 9). The mechanism also predicts which margin carries the response—because NFTs are indivisible, the volume response operates primarily through the transaction-count (extensive) margin rather than through per-trade prices—and whether information content matters more than attention; the next two subsections test these predictions. Section 4 formalizes the comparative statics.

3.4 Mechanism: Transactions vs. Prices

The wedge mechanism makes a sharper prediction about *which* margin responds to communication. If fees suppress volume by keeping marginal investors out of the market, the response should appear in transaction counts rather than per-trade prices. Since Volume = Transactions × Avg Price, in logs:

$$\log(\text{Volume}) = \log(\text{Transactions}) + \log(\text{Avg Price}).$$

We estimate the coefficient on log(Tweets+1) separately for each component. Note that this decomposition necessarily restricts to collection-weeks with positive transactions ($N = 9{,}928$), since average prices are undefined for zero-trade weeks. The sample conditioning drops the zero-trade observations where the no-trade mechanism is most directly operative; the estimates therefore capture the intensive-margin decomposition conditional on positive trading activity. The regressor is $\log(\text{Tweets} + 1)$ rather than the raw tweet count used in the baseline semi-elasticity, so coefficient levels are not directly comparable to Tables 3–5. Table 36 in the Appendix presents a two-part (hurdle) decomposition that explicitly accounts for zero-trade weeks: a linear probability model for whether a collection trades in a given week, followed by the intensive-margin OLS conditional on positive activity.

Table 11: Channel Decomposition: Transactions vs. Prices (Ethereum)

	(1) Log(Transactions)	(2) Log(Avg Price)
Log(Tweets + 1)	0.190***	0.090**
	(0.044)	(0.039)
Observations	9,928	9,928
Collection FE	Yes	Yes
Week FE	Yes	Yes

Notes: Estimates from equation (1) with collection and week fixed effects, restricting to collection-weeks with positive transactions (so that average prices are defined). Standard errors are two-way clustered by collection and week.

Table 11 and Figure 6 show that tweets load more heavily on transaction counts than on prices. The transaction elasticity is roughly twice the price elasticity (0.190 vs. 0.090), so about two-thirds of the volume response operates through the extensive margin—additional trades rather than per-trade price changes.

Breaking down the channel decomposition by enforcement period (Table 12), the volume semi-elasticity is 1.28% in the pre-erosion period, falls to 0.37% during partial erosion, and remains at 0.46% under full erosion. The transaction semi-elasticity follows a similar U-shaped pattern (1.04% → 0.30% → 0.48%), while the price semi-elasticity is small and insignificant throughout (0.27% → 0.05% → 0.08%). The attenuation concentrates on the extensive margin.

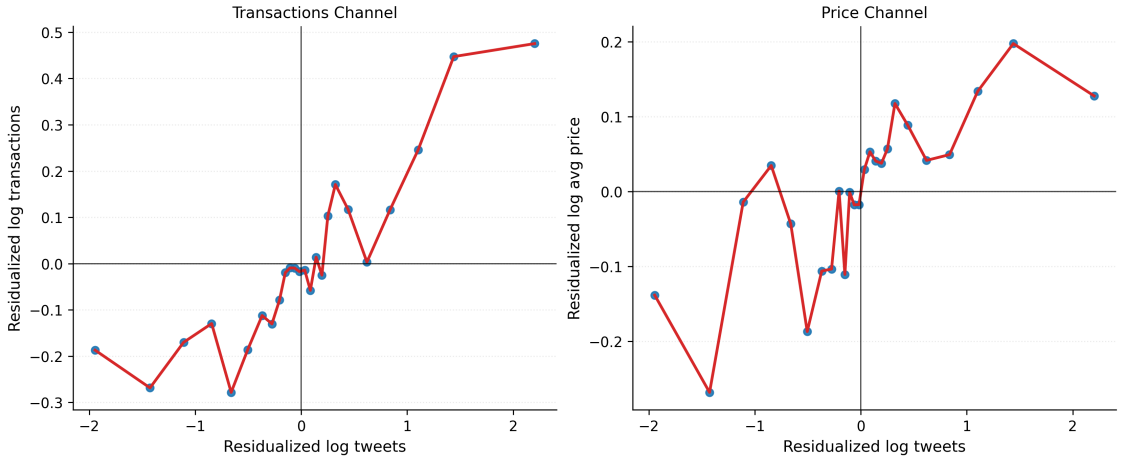

Figure 6: Channel Decomposition: Tweets Are More Strongly Associated with Transactions Than Prices

Notes: Binned scatter plots of residualized log outcomes against residualized log tweets, after partialling out collection and week fixed effects. The sample restricts to collection-weeks with positive transactions (so that average prices are defined). Left panel: log transactions. Right panel: log average price. The steeper slope for transactions indicates that tweets are primarily associated with trading activity rather than per-trade prices.

Table 12: Channel Decomposition Across Enforcement Regimes (Ethereum)

	(1) Log(Volume)	(2) Log(Transactions)	(3) Log(Avg Price)
Tweet count × Pre-erosion	0.0128**	0.0104***	0.0027
	(0.0050)	(0.0031)	(0.0030)
Tweet count × Partial erosion	0.0037	0.0030*	0.0005
	(0.0029)	(0.0018)	(0.0012)
Tweet count × Full erosion	0.0046**	0.0048***	0.0008
	(0.0022)	(0.0016)	(0.0021)
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations		9,928	

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with three regime indicators for each outcome. Coefficients represent percentage change per additional tweet. Sample restricted to collection-weeks with positive transactions (N=9,928), so coefficient levels differ from the full-sample baseline estimates. Wald tests report p-values for pre-erosion vs. full-erosion equality within each outcome. Standard errors are two-way clustered by collection and week.

3.4.1 Information vs. Attention: Announcements

A tweet can move volume because it carries information or because it draws attention. We separate the two by classifying tweets into *announcements* (roadmap updates, partnerships, airdrops) versus other communication. Classification relies on a zero-shot language model; misclassification would attenuate the estimated gap between announcement and non-announcement tweets, so the true difference may be larger.

Announcements predict trading far more strongly than other tweets (Table 13): one additional announcement tweet is associated with a 6.05% increase in volume, while one additional non-announcement tweet is associated with only a 0.35% increase ($p = 0.002$ for the difference)—a 17-fold difference in effect size. Announcements load on both transaction counts and average prices; non-announcement tweets show up mainly in transactions, with little link to prices. The asymmetry points toward information rather than pure attention, though the two channels are not fully separable.

Table 13: Information vs. Attention in Creator Communication (Ethereum)

	Log(Volume)
Announcement tweet count	0.0605***
	(0.0187)
Non-announcement tweet count	0.0035**
	(0.0017)
p-value: Ann vs. Non-ann	0.002
Collection FE	Yes
Week FE	Yes
Observations	17,403

Notes: Semi-elasticity specification: announcement tweet count and non-announcement tweet count (not log-transformed) enter separately. Coefficients represent percentage change in volume per additional tweet of each type. Tweets are classified using a zero-shot language model; misclassification would attenuate the estimated gap. Standard errors are two-way clustered by collection and week. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Figure 7 displays the contrast.

Interacting both tweet types with the weak-enforcement indicator sharpens the picture (Table 14, Figure 8). The announcement semi-elasticity drops from 6.67% pre-policy to 4.69% post-policy (interaction: −1.98 pp); the semi-elasticity for non-announcement tweets remains essentially flat (0.42% to 0.33%, interaction: −0.09 pp). The within-regression difference (Announcements × Post) − (Other × Post) is negative (−0.019), though the individual interaction coefficients are imprecisely estimated and the significance of the difference depends on the off-diagonal covariance. The overall attenuation in Table 5 is thus driven by information-rich communication, not by a uniform decline.

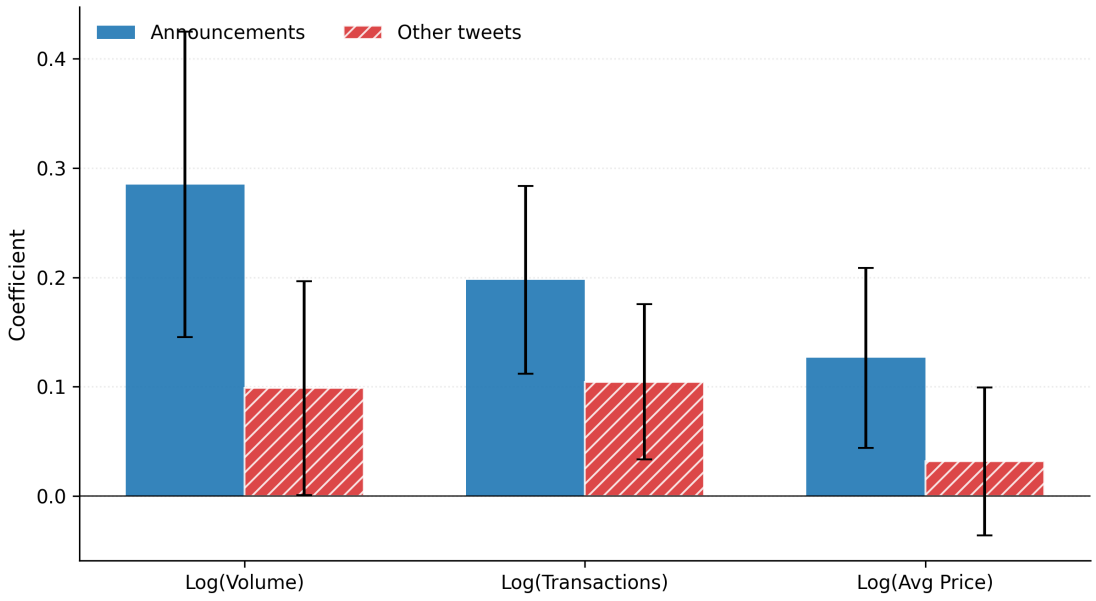

Figure 7: Announcements Are More Strongly Associated with Trades and Prices Than Other Tweets (Ethereum)

Notes: Coefficients and 95% confidence intervals from Table 13. Announcements are tweets classified as providing concrete project updates; other tweets include general communication and promotion.

Table 14: Information vs. Attention Around the Fee-Regime Change (Ethereum)

	Log(Volume)
Announcement tweet count	0.0667***
	(0.0240)
Non-announcement tweet count	0.0042
	(0.0032)
Announcement × Post-policy	-0.0198
	(0.0269)
Non-announcement × Post-policy	-0.0009
	(0.0037)
Collection FE	Yes
Week FE	Yes
Observations	17,403

Notes: Semi-elasticity specification: announcement and non-announcement tweet counts (not log-transformed) are each interacted with post-policy indicator. Coefficients represent percentage change in volume per tweet. Interaction coefficients show how the semi-elasticity changes post-policy. Standard errors are two-way clustered by collection and week. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

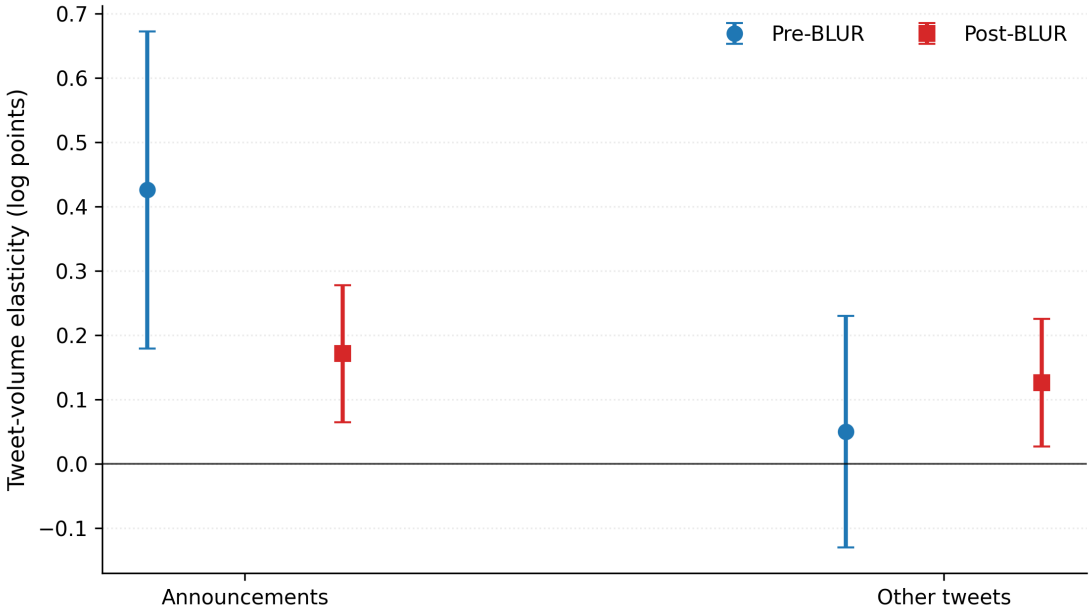

Figure 8: The Announcement-Volume Link Declines After the Fee-Regime Change (Ethereum)

Notes: Coefficients and 95% confidence intervals implied by Table 14. Standard errors are two-way clustered by collection and week.

The announcement decomposition across enforcement periods (Table 15) shows the announcement semi-elasticity is 6.66% pre-erosion, drops to 3.75% during partial erosion, and remains at 4.93% under full erosion. Non-announcement tweets show a flat pattern across regimes (0.42%, 0.32%, 0.33%) (Figure 9).

A text-based informativeness score confirms the pattern (Table 16). The negative main effect on tweet count reflects evaluation at mean informativeness and pre-policy baseline; the economic effect depends on the interaction terms. Because this regression restricts to tweet-active weeks and interacts tweet count with demeaned informativeness, the main coefficient is not directly comparable to the baseline semi-elasticity. The interaction coefficient between tweet count and informativeness is 0.31 pre-policy. Since informativeness is demeaned but not standardized, a one-standard-deviation increase (SD = 0.044; Table 2) raises the semi-elasticity by approximately $0.31 \times 0.044 \approx 1.4$ percentage points. Post-policy, this interaction coefficient falls by 0.19 (triple interaction: -0.19, $p = 0.01$), indicating weakened sensitivity to content quality. An Online Appendix decomposition pins the attenuation to the transactions margin—the conditional price response does not decline.

Splitting by enforcement period (Table 17), the raw interaction coefficient between tweet count and informativeness is 0.31 pre-erosion, falls to 0.12 during partial erosion, and remains at 0.11 under full erosion. Since informativeness is demeaned with SD = 0.044, these raw coefficients translate to one-SD effects of approximately 1.4, 0.5, and 0.5 percentage points respectively—attenuated sensitivity to tweet informativeness that remains positive but becomes statistically

Table 15: Information vs. Attention Across Enforcement Regimes (Ethereum)

	Pre-erosion	Partial erosion	Full erosion
Announcement tweet count	0.0666***	0.0375	0.0493***
	(0.0241)	(0.0243)	(0.0176)
Non-announcement tweet count	0.0042	0.0032	0.0033*
	(0.0032)	(0.0030)	(0.0017)
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations		17,403	

Notes: Semi-elasticity specification: announcement tweet count and non-announcement tweet count (not log-transformed) are each interacted with three regime indicators. Coefficients represent percentage change in volume per tweet. Tweets are decomposed using weekly announcement share (soft assignment). Wald tests compare pre-erosion vs. full-erosion coefficients. Standard errors are two-way clustered by collection and week.

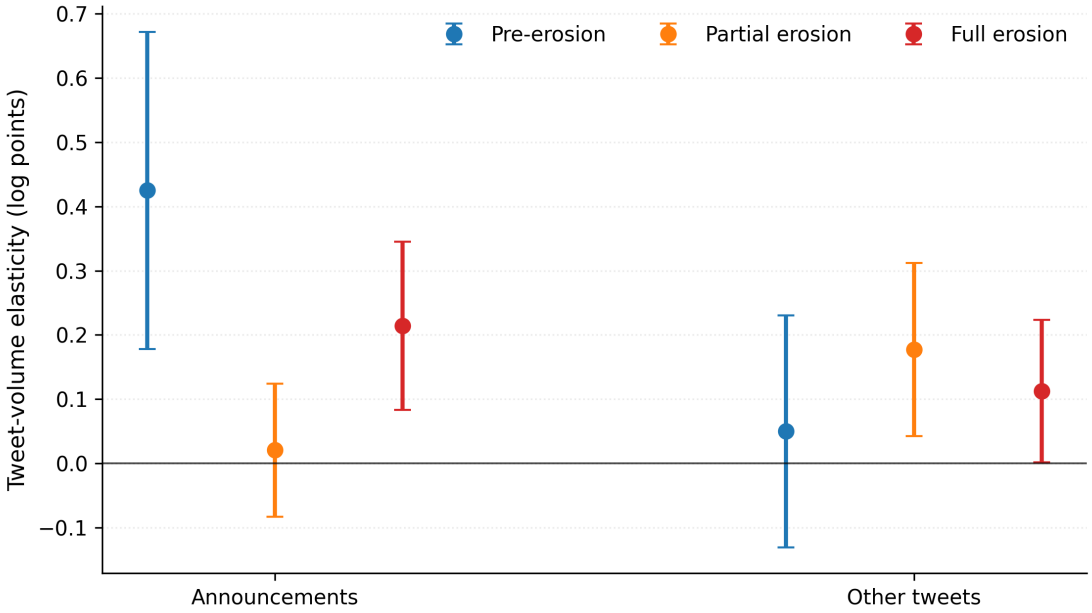

Figure 9: Announcement and Other-Tweet Semi-Elasticities Across Enforcement Regimes (Ethereum)

Notes: Point estimates and 95% confidence intervals from Table 15 for each enforcement regime.

Table 16: Tweet Informativeness and Trading Around the Fee-Regime Change (Ethereum)

	Log(Volume)
Tweet count	-0.1070***
	(0.0234)
Tweet count × Informativeness	0.3107***
	(0.0682)
Tweet count × Post-policy	0.0699***
	(0.0268)
Tweet count × Informativeness × Post-policy	-0.1919**
	(0.0774)
Collection FE	Yes
Week FE	Yes
Observations	17,403

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with informativeness (mean composite score, demeaned) and post-policy indicator. The main effect on tweet count corresponds to informativeness evaluated at its sample mean. Interaction coefficients show how informativeness moderates the semi-elasticity, and how this moderation changes post-policy. Sample restricted to collection-weeks with at least one tweet. Standard errors are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

imprecise. The transaction-count margin drives the decline; the price margin follows a similar but noisier path.

3.5 Cross-Blockchain Evidence: Solana

If the Ethereum attenuation reflects a fee-wedge channel rather than an Ethereum-specific shock, a similar pattern should appear on Solana, where Magic Eden weakened royalty enforcement in October 2022. Solana differs from Ethereum in architecture, gas costs (near zero), market composition, and collection characteristics. However, FTX and Alameda were deeply embedded in the Solana ecosystem—SOL fell more than 50% at the FTX collapse, far exceeding the 15–25% decline in other major blockchain tokens (e.g., Ethereum, Avalanche)[15]—so the FTX shock is not orthogonal across chains. A qualitative replication of the *fee-heterogeneity* pattern is the more diagnostic test, since wash trading and aggregate shocks do not predict differential attenuation by royalty rate.

Pre-erosion semi-elasticities are substantially larger on Solana (0.0252, or 2.52% per tweet) than on Ethereum (0.0103, or 1.03% per tweet), reflecting differences in market design and participant composition. The decline after enforcement weakened—from 2.52% to 0.72% (= 2.52% - 1.79%), a 71% reduction—matches the Ethereum pattern in sign and is economically large, though

[15] "Solana's Close Ties to FTX Cause Financial and Structural Damage," *Protos*, November 2022.

Table 17: Tweet Informativeness Across Enforcement Regimes (Ethereum)

	Pre-erosion	Partial erosion	Full erosion
Tweet count	-0.1069***	-0.0388	-0.0348
	(0.0235)	(0.0391)	(0.0240)
Tweet count × Informativeness	0.3106***	0.1222	0.1132
	(0.0685)	(0.1113)	(0.0690)
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations		17,403	

Notes: Semi-elasticity specification: tweet count (not log-transformed) is interacted with informativeness (mean composite score, demeaned) in each regime. Coefficients represent interaction effects showing how informativeness moderates the tweet-volume semi-elasticity. Sample restricted to collection-weeks with at least one tweet. Standard errors are two-way clustered by collection and week.

Table 18: The Tweet-Volume Relationship Before and After the Policy Change (Solana)

	(1)	(2)
Tweet count	0.0252***	0.0071***
	(0.0077)	(0.0023)
Tweet count × Post-policy	-0.0179	-0.0055**
	(0.0110)	(0.0028)
Collection FE	Yes	Yes
Week FE	Yes	Yes
Lagged volume	No	Yes
Observations	12,450	12,367

Notes: Same specification as Table 5 applied to the Solana sample. The pre-erosion semi-elasticity (2.52%) is substantially larger than on Ethereum (1.03%), and the post-erosion decline is also more pronounced. Standard errors are two-way clustered by collection and week. Appendix Table 32 reports the log–log specification. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

imprecisely estimated. For announcement decomposition on Solana, we use the log–log specification due to lower tweet counts per collection-week; the announcement elasticity is 0.64 pre-policy and declines to 0.22 post-policy (Appendix Table 19).

Appendix Table 21 and Figure 26 decompose the Solana association into transactions and prices. Unlike Ethereum, where transactions account for roughly two-thirds of the volume response (Table 11), on Solana the price channel is comparable in magnitude. With near-zero gas fees, fewer Solana investors sit at the participation hurdle, and more of the adjustment occurs through prices rather than additional entries. The model has a single proportional cost r and does not distinguish gas from royalties. Formally predicting the Solana pattern would require a richer cost structure that separates the two components; the qualitative intuition—lower total wedge means fewer marginal investors—is consistent with the model's comparative statics.

The appendix replicates each remaining Ethereum test on Solana: announcements vs. other tweets (Table 19), informativeness (Table 20), and fee heterogeneity (Table 22). The key qualitative result—that the tweet-volume relationship attenuates after the policy change—replicates on both chains, though the coefficient signs on the fee-level interactions in Table 22 differ from Ethereum (Table 6), likely reflecting differences in functional form (semi-elasticity on Ethereum vs. log-log on Solana) and sample composition. Precision is lower in the smaller sample. On Solana the announcement-specific pattern is noisier and not statistically distinguishable from zero, reflecting both the smaller sample and different market composition. Cross-chain level differences reflect market structure, investor composition, and network fees; the relevant comparison is within each blockchain, where both show attenuation after enforcement weakened.

Appendix Figure 22 provides the Solana-only event study; Appendix Figure 23 plots partial correlations between tweet informativeness and trading volume on both blockchains.

3.6 Reverse Causality

The evidence so far is contemporaneous. The natural concern is reverse causality: creators tweet *in response* to trading activity, and the estimated association overstates how much communication precedes trading.

We separate the two directions with local projections (Jordà, 2005). For horizons $h = 0, \ldots, 8$ weeks:

$$y_{i,t+h} = \alpha_{i,h} + \delta_{t,h} + \beta_h \, x_{i,t} + \sum_{p=1}^{P} \phi_{h,p} \, y_{i,t-p} + \sum_{p=1}^{P} \psi_{h,p} \, x_{i,t-p} + \varepsilon_{i,t+h}, \tag{4}$$

Two lags of both variables absorb predictable dynamics; β_h captures the marginal predictive content of $x_{i,t}$ for the outcome h weeks ahead. We run the regression in both directions: *forward* ($y = \log(1+\text{Volume})$, $x = \log(1+\text{Tweets})$) and *reverse* (roles swapped).

Ethereum. Figure 10 plots both paths. In the forward direction, tweets predict volume at all horizons: $\hat{\beta}_0 = 0.141$ ($p < 0.001$), fading to $\hat{\beta}_8 = 0.060$ ($p = 0.014$). Persistence out to $h = 8$ is consistent with gradual diffusion across a heterogeneous investor base. In the reverse direction, coefficients are roughly a third as large ($\hat{\beta}_0 = 0.046$) and fluctuate between marginal and moderate significance. Volume predicts tweets to some degree—creators respond to market events—but the reverse channel is weaker at every horizon.

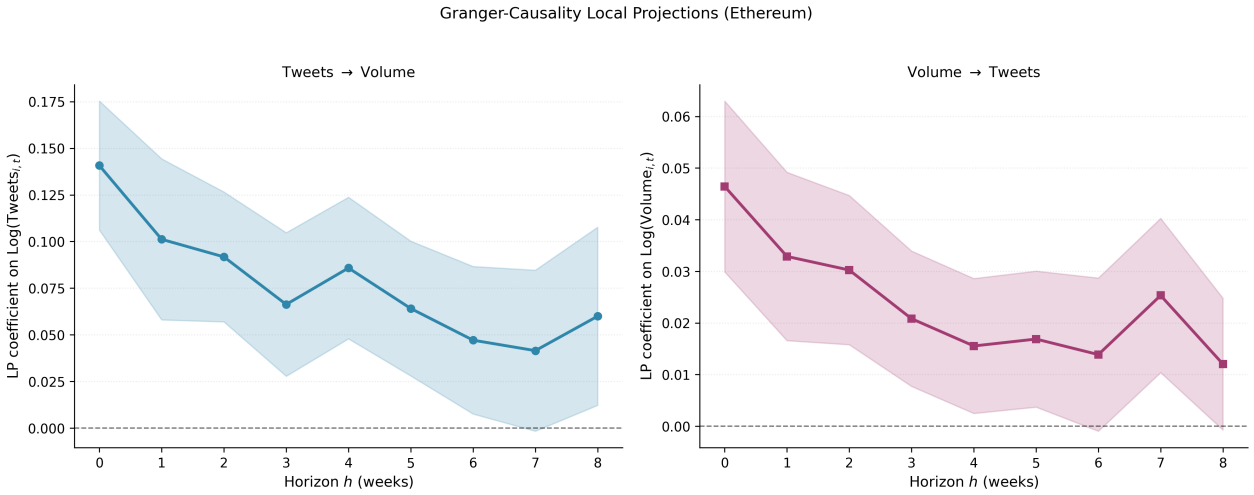

Figure 10: Local Projections: Tweets and Volume (Ethereum)

Notes: Coefficients $\hat{\beta}_h$ from equation (4) with $P = 2$ lags of both variables. Left panel: $y = \log(1+\text{Volume})$, $x = \log(1+\text{Tweets})$ (forward direction). Right panel: roles reversed (reverse direction). Shaded bands are pointwise 95% confidence intervals based on standard errors two-way clustered by collection and week.

Solana. The asymmetry is sharper on Solana (Figure 11). The forward path is significant at every horizon and *increases* from $\hat{\beta}_0 = 0.17$ to $\hat{\beta}_8 = 0.33$, consistent with a cumulative response. Reverse coefficients are an order of magnitude smaller ($\hat{\beta}_0 = 0.018$) and lose significance from $h = 4$ onward. Volume exhibits limited feedback into communication beyond a few weeks.

On both chains the predictive relationship runs predominantly from communication to trading. Some feedback exists—particularly on Ethereum—but the forward channel dominates at every horizon. Appendix Tables 33 and 34 report all coefficients.

Hourly data provide sharper evidence on temporal ordering. Using individual tweet timestamps and transaction records at hourly frequency, we conduct three tests (Online Appendix). First, an event study around 92,891 Ethereum tweets (41,867 on Solana) shows abnormal volume flat in the 24 hours before a tweet, jumping at $h = 0$ and staying elevated for 48 hours (Figure 12). Second, a mirror test centered on volume spikes (hours above the collection-specific 90th percentile) shows that tweet activity does *not* rise before trading surges. Third, the hourly cross-correlation function peaks at lag $k = +1$ on Ethereum (CCF $= 0.038$) and at $k = 0$ on Solana, with positive lags

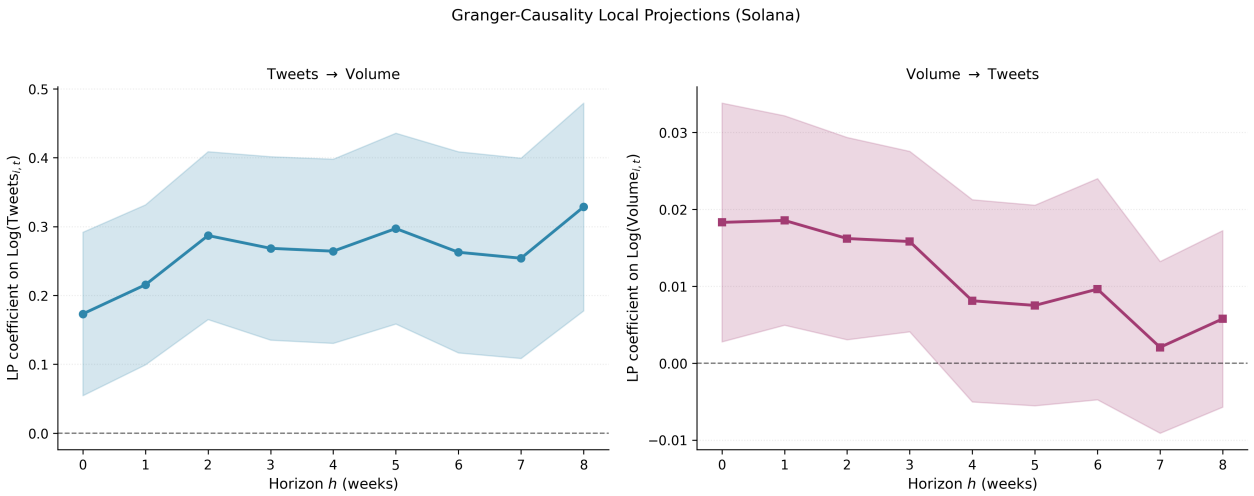

Figure 11: Local Projections: Tweets and Volume (Solana)

Notes: Same specification as Figure 10 applied to the Solana sample ($N = 12{,}450$). Standard errors are two-way clustered by collection and week.

systematically exceeding negative lags. These patterns are more consistent with tweets preceding volume than the reverse.

3.7 Robustness

Appendix C presents robustness checks: FTX-period exclusion, winsorization, first differences, alternative clustering, and a Poisson PML estimator that handles the $\log(1+x)$ zeros directly (Appendix C.1).

The association is stable across market regimes. Splitting by ETH price terciles (Appendix Table 24), the semi-elasticity is positive and significant in all three, strongest in high-price periods. A placebo exercise (Appendix Table 23, Figure 28) finds the observed slope break around the actual policy dates is indistinguishable from random variation (randomization p=0.93 for Ethereum, p=0.49 for Solana). This is consistent with the absence of a discrete structural break at any single policy date, as expected under progressive erosion in the semi-elasticity as enforcement weakened across multiple marketplace decisions. Figure 29 shows approximate linearity in the residualized relationship. The Online Appendix reports time-series diagnostics (unit roots, cointegration, Granger tests in levels).

4 Model

The empirical findings motivate the model's key features: volume responds to communication through transaction counts rather than prices; proportional fees and congestion amplify the response; both relationships weaken when fees decline. We embed proportional creator fees in a noisy

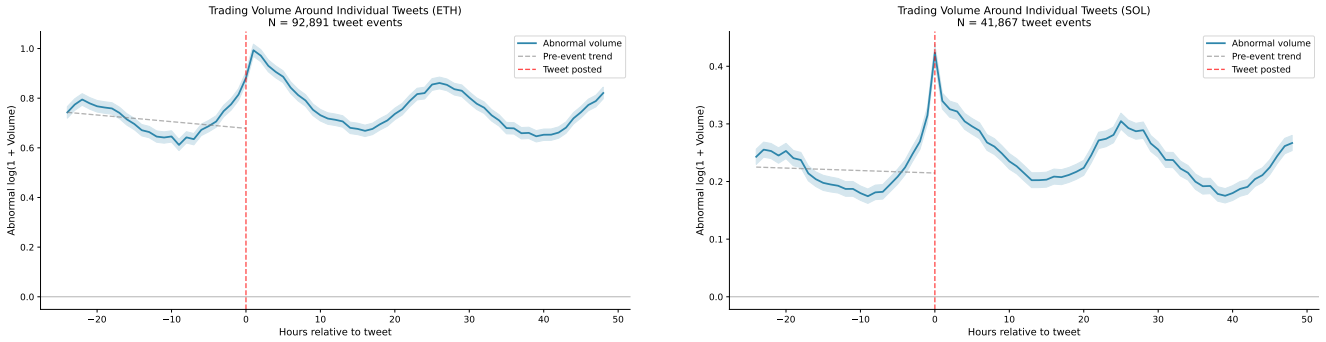

Figure 12: Trading volume around individual tweets at hourly frequency. Abnormal log(1 + volume) (relative to a collection-specific mean in the pre-event window) is flat in the 24 hours before a tweet, spikes at $h = 0$, and remains elevated for roughly 48 hours. The pattern holds on both Ethereum (92,891 tweets, left) and Solana (41,867 tweets, right). Shaded bands are 95% confidence intervals. The trend line is a local linear fit to the pre-tweet window.

rational expectations equilibrium augmented with *heterogeneous interpretation* of public signals to account for these patterns, derive predictions about inaction regions and disclosure timing, and show that welfare improvements come primarily from eliminating the per-trade distortion, with an additional channel operating through changing creators' information-disclosure incentives under early-resolution-of-uncertainty preferences.

The key innovation relative to the standard CARA–Normal noisy REE is that the public signal is not commonly observed. Instead, each investor sees the creator's public communication with idiosyncratic interpretation noise, following the spirit of Harris and Raviv (1993) and Kandel and Pearson (1995). When interpretation precision is finite, increased public precision *widens* cross-sectional belief dispersion rather than compressing it, reversing the standard comparative static and matching the empirical finding that creator tweets raise trading volume.

Consider a two-period economy ($t \in \{1, 2\}$) in which a risky asset with unit net supply trades competitively each period; liquidation occurs after $t = 2$ trading. The asset represents an NFT collection's residual value, and a founder chooses how much public information to release each period while collecting proportional fees on trades. The framework extends Admati (1985) to incorporate proportional fees and endogenous disclosure.

A unit mass of competitive investors have CARA preferences with common risk aversion $H > 0$:

$$U_i(W_{i,2}) = -\exp(-HW_{i,2}). \tag{5}$$

Aggregate net supply equals one, absorbed by the investor continuum and noise traders. Vives (2008) provides a textbook treatment of the CARA-Gaussian framework.

Assumption 4.1 (Initial holdings)**.** *All investors begin with zero holdings:* $x_{i,0} = 0$ *for all* $i \in [0, 1]$.

Zero initial holdings ensure that date-1 trading reflects information-driven demand rather than

rebalancing—a natural assumption for newly issued NFTs.

The asset's terminal payoff f is drawn from a Normal prior:

$$f \sim \mathcal{N}(\theta, \sigma_f^2). \tag{6}$$

Information arrives through two channels. Each investor i observes a private signal

$$s_{i,t} = f + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim \mathcal{N}(0, \tau_t^{-1}), \tag{7}$$

where $\tau_t > 0$ denotes private signal precision. The founder controls a public signal

$$y_t = f + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \kappa_t^{-1}), \tag{8}$$

with precision $\kappa_t \geq 0$ chosen subject to an aggregate disclosure budget:

$$\kappa_1 + \kappa_2 = \bar{\kappa}, \qquad \bar{\kappa} \geq 0. \tag{9}$$

Heterogeneous interpretation. The single departure from the standard framework is that investor i does not observe y_t directly. Instead, each investor observes

$$\hat{y}_{i,t} = y_t + \xi_{i,t} = f + \eta_t + \xi_{i,t},$$

where $\xi_{i,t} \sim \mathcal{N}(0, \psi_t^{-1})$ is an idiosyncratic interpretation error, independent across investors, periods, and of all other random variables. The parameter $\psi_t > 0$ is the *interpretation precision*. Since $\hat{y}_{i,t}$ is investor-specific, it functions as an additional private signal about f with effective precision

$$\kappa_t^{\text{eff}} \equiv \frac{\kappa_t \, \psi_t}{\kappa_t + \psi_t},$$

which is the harmonic mean of κ_t and ψ_t. When $\psi_t \to \infty$, $\kappa_t^{\text{eff}} \to \kappa_t$ and the public signal is commonly observed, recovering the standard model.

All random variables are mutually independent across i, t, and signal types.

Noise traders generate order flow $u_t \sim \mathcal{N}(0, \sigma_{u,t}^2)$, independent across periods. Market clearing requires

$$\int_0^1 x_{i,t} \, di + u_t = 1. \tag{10}$$

A competitive market-clearing price P_t equates aggregate demand to fixed supply at each date, abstracting from the strategic considerations in Kyle (1985) to isolate fee structure's effects on trading behavior.

Trading is subject to a proportional fee (royalty) $r \in [0, 1)$ collected by the founder:

$$\text{Fee}_{i,t} = rP_t|x_{i,t} - x_{i,t-1}|. \tag{11}$$

This fee structure—standard in NFT marketplaces—creates a wedge between buying and selling prices that scales with the current market price.

Within each period the timing is: (i) the founder commits to public precision κ_t; (ii) the public signal y_t realizes; (iii) each investor observes a private signal $s_{i,t}$ and an interpreted public signal $\hat{y}_{i,t}$; (iv) investors submit demand schedules; (v) the market clears at price P_t; (vi) at $t = 2$, the payoff f realizes. Let $I_{i,t} = \sigma(\theta, \hat{y}_{i,\leq t}, s_{i,\leq t}, P_{\leq t})$ denote investor i's information set at date t. Posterior beliefs incorporate the interpreted public signals:

$$f \mid \hat{y}_{i,\leq t}, s_{i,\leq t} \sim \mathcal{N}\big(\mu_{i,t},\ \sigma^2_{i,t}\big)\,. \tag{12}$$

4.1 Equilibrium Prices and Information Aggregation

We conjecture a linear price function, verify that it clears the market, and show that the resulting coefficients have closed forms and the linear REE is unique.

Proportional trading fees kink investors' demands; aggregate demand need not be perfectly linear. We therefore characterize the linear REE in the frictionless benchmark ($r = 0$), which pins down how prices load on fundamentals and noise, and then study trading *with fees* conditional on those prices. This delivers the inaction regions and volume expressions used to interpret the empirical "information × transaction-cost" interactions.

Equilibrium prices take the form

$$P_t = \alpha_t + \beta_t f + \sum_{s=1}^{t} \lambda_{t,s}\eta_s + \gamma_t u_t, \tag{13}$$

where $(\alpha_t, \beta_t, \lambda_{t,1}, \ldots, \lambda_{t,t}, \gamma_t)$ are constants to be determined. Unlike the standard noisy REE, posterior updating below omits any additional price-learning term: prices clear the market but are not treated as a separate signal about fundamentals.

Bayesian updating yields each investor's posterior. Investor i's information set at date t includes all interpreted public signals, all private signals, and all past prices: $I_{i,t} = \sigma(\theta, \hat{y}_{i,\leq t}, s_{i,\leq t}, P_{\leq t})$. In the Gaussian conjugate framework, the posterior precision accumulates all observed signal precisions. At date t, define the *cumulative* posterior precision:

$$\sigma^{-2}_{i,t} = \sigma^{-2}_f + \sum_{s=1}^{t} (\kappa^{\text{eff}}_s + \tau_s), \tag{14}$$

which incorporates both interpreted public signals (effective precisions κ^{eff}_s) and private signals

(precisions τ_s) from all periods through t. Since τ_s and κ_s^{eff} are common across investors, $\sigma_{i,t}^2$ does not vary across investors. The price does not enter as an additional signal (Assumption 4.3 below).

The posterior mean at date t is:

$$\mu_{i,t} = \sigma_{i,t}^2 \left[\sigma_f^{-2}\theta + \sum_{s=1}^{t} \left(\tau_s \, s_{i,s} + \kappa_s^{\text{eff}} \hat{y}_{i,s} \right) \right]. \tag{15}$$

Investors differ in posterior means through their private signals *and* their idiosyncratic interpretation of public signals.

Remark 4.2 (Cumulative information and disclosure comparative statics)**.** *At date 2, the full posterior precision is*

$$\sigma_{i,2}^{-2} = \sigma_f^{-2} + \kappa_1^{\text{eff}} + \kappa_2^{\text{eff}} + \tau_1 + \tau_2.$$

When ψ_t is the same in both periods, $\kappa_t^{\text{eff}} = \kappa_t\psi/(\kappa_t + \psi)$ is a concave function of κ_t. Unlike the standard model where $\sigma_{i,2}^{-2}$ depends only on total disclosure $\bar{\kappa}$, here the effective precisions $\kappa_1^{\text{eff}} + \kappa_2^{\text{eff}}$ can depend on the split (κ_1, κ_2). For fixed interpretation precision ψ, concavity implies that $\kappa_1^{\text{eff}} + \kappa_2^{\text{eff}}$ is maximized by smoother disclosure and minimized by concentrating all disclosure in one period. The corner disclosure results below therefore come from fee structure and trading incentives, not from a mechanical gain in total effective precision. Under royalties, the founder maximizes volume rather than precision. The volume function's dependence on the threshold structure can create a countervailing incentive to concentrate disclosure; the net direction depends on the relative strength of the precision-loss and disagreement-widening effects.

In a frictionless benchmark ($r = 0$), CARA-Normal preferences imply myopic demand

$$x_{i,t}^{(0)} = \frac{\mu_{i,t} - P_t}{H\sigma_{i,t}^2}. \tag{16}$$

Averaging the posterior mean across investors and imposing market clearing pins down the price coefficients.

Assumption 4.3 (No additional price learning)**.** *In the heterogeneous-interpretation model, posterior updating does not include an additional price-learning term:*

$$\phi_{P,t} = 0. \tag{17}$$

This restriction shuts down the standard REE channel through which investors infer additional information from equilibrium prices. In the present draft, disagreement is driven entirely by private signals and heterogeneous interpretation of creator communication. Prices still clear the market and depend on fundamentals, realized public-signal innovations, and noise-trader demand, but they do not feed back into posterior precision as a separate signal. Relaxing this assumption to

allow partial price learning ($\phi_{P,t} > 0$) would add a common-signal component that partially offsets the disagreement channel. The qualitative results survive as long as effective interpretation-noise precision exceeds price-signal precision, so that the widening effect dominates. Market clearing (10) implies

$$\bar{\mu}_t - P_t = H\sigma_{i,t}^2(1 - u_t), \tag{18}$$

where $\bar{\mu}_t := \int_0^1 \mu_{i,t}\, di$ is the average posterior mean.

Under Assumption 4.3, the remaining coefficients solve a triangular system. Matching coefficients on f, on the public-signal innovations η_s, and on u_t in market clearing gives:

$$\beta_t = \frac{\sum_{s=1}^{t}(\tau_s + \kappa_s^{\text{eff}})}{\sigma_{i,t}^{-2}} = 1 - \sigma_f^{-2}\sigma_{i,t}^2, \tag{19}$$

$$\lambda_{t,s} = \frac{\kappa_s^{\text{eff}}}{\sigma_{i,t}^{-2}} = \sigma_{i,t}^2\kappa_s^{\text{eff}}, \qquad s = 1, \ldots, t, \tag{20}$$

$$\gamma_t = H\sigma_{i,t}^2. \tag{21}$$

The price loading β_t on fundamentals increases in the cumulative signal precision $\sum_{s=1}^{t}(\tau_s + \kappa_s^{\text{eff}})$ and decreases in prior precision σ_f^{-2}. The coefficients $\lambda_{t,s}$ load the realized creator-signal innovations into prices because those innovations shift average posterior beliefs in the investor continuum. The noise loading γ_t is proportional to risk aversion H.

The intercept α_t determines the expected price level and enters the welfare comparison below.

Proposition 4.4 (Price intercept and expected price)**.** *The equilibrium intercept satisfies*

$$\alpha_t = (1 - \beta_t)\theta - H\sigma_{i,t}^2, \tag{22}$$

and the unconditional expected price is

$$\mathbb{E}[P_t] = \theta - H\sigma_{i,t}^2. \tag{23}$$

Proof. Averaging the demand function across investors and matching the constant term in market clearing gives $\alpha_t = \sigma_{i,t}^2 \cdot \sigma_f^{-2}\theta - H\sigma_{i,t}^2$. Since $1 - \beta_t = \sigma_{i,t}^2 \cdot \sigma_f^{-2}$, this is equivalent to (22). Taking unconditional expectations of the price function then gives $\mathbb{E}[P_t] = \alpha_t + \beta_t\theta = \theta - H\sigma_{i,t}^2$. □

Remark 4.5 (Full posterior versus price-only variance)**.** *The risk discount $H\sigma_{i,t}^2$ incorporates all sources of residual uncertainty, including the information content of private signals and interpreted public signals. The expected price increases in total posterior precision.*

Uniqueness follows from the triangular structure.

Theorem 4.6 (Existence and uniqueness)**.** *For any $(\sigma_f^{-2}, \tau_t, \kappa_t^{\text{eff}}, H, \sigma_{u,t}^2)$ with $\tau_t > 0$, $H > 0$, and $\sigma_{u,t}^2 > 0$, there exists a unique set of equilibrium price coefficients.*

Proof. Under Assumption 4.3, the posterior precision $\sigma_{i,t}^{-2}$ is pinned by primitives. Coefficient matching yields a triangular system with strictly positive diagonal entries, so the solution is unique. □

The triangular structure eliminates fixed-point arguments or contraction mappings that arise in more general noisy REE models—with no price learning, posterior precision is pinned by primitives, so coefficients solve sequentially.

Cross-sectional distribution of mispricing. A central object is the cross-sectional distribution of investor-level mispricings, which determines who trades. Under heterogeneous interpretation, conditional on the common states $(f, \eta_{\leq t}, u_t)$, the mispricing $Z_{i,t} := \mu_{i,t} - P_t$ is Normally distributed across investors with cross-sectional variance

$$\sigma_{Z,t}^2 = \frac{1}{(\sigma_{i,t}^{-2})^2} \left(\sum_{s=1}^{t} \tau_s + \sum_{s=1}^{t} \frac{(\kappa_s^{\text{eff}})^2}{\psi_s} \right). \tag{24}$$

The first sum is the standard private-signal dispersion, present in both models. The second sum is the *disagreement component*, present only when $\psi_s < \infty$. It captures cross-sectional dispersion due to heterogeneous interpretation of creator communication. At date 1, (24) collapses to the one-period expression used in the calibration figures. Appendix D derives (24) from the full posterior and shows that current disclosure raises $\sigma_{Z,t}$ once current effective precision exceeds an explicit threshold determined by inherited disagreement and cumulative precision.

In the standard REE ($\psi_t = \infty$), the disagreement component vanishes, $\sigma_{Z,t}^2 = \left(\sum_{s=1}^{t} \tau_s\right) / (\sigma_{i,t}^{-2})^2$ is *strictly decreasing* in cumulative public precision, and more public information compresses belief dispersion. At the baseline calibration ($\sigma_f = 1$, $\tau = 1$, $\psi = 2$, $\bar{\kappa} = 2$), the threshold condition (53) is satisfied for $\kappa_t \geq 0.75$ at date 1, well within the empirically relevant range.

4.2 Trading Behavior and Inaction Regions

Proportional fees create inaction regions where investors refrain from trading despite holding informative signals. We characterize optimal trading at each date, working backward from the terminal period. The no-trade band arises because small gains from rebalancing cannot justify the fee (Constantinides, 1986).

Remark 4.7 (Price formation with fees)**.** *We treat equilibrium prices as the frictionless rational expectations outcomes derived in the previous subsection ($r = 0$), and analyze individual trading decisions conditional on these prices when $r > 0$. This is a partial-equilibrium approximation: proportional fees enter individual optimization but do not feed back into market clearing, price coefficients, or price informativeness. The approach simplifies the exposition and delivers the qualitative predictions (inaction regions, extensive-margin response, congestion interactions) that*

map to the data; normative conclusions about welfare and disclosure timing should be understood as conditional on this approximation.

At date 2, investor i enters with position $x_{i,1}$ and solves

$$\max_{x_{i,2}\in\mathbb{R}} \quad x_{i,2}(\mu_{i,2} - P_2) - \frac{H}{2}\sigma_{i,2}^2 x_{i,2}^2 - rP_2|x_{i,2} - x_{i,1}|. \tag{25}$$

The objective is strictly concave in $x_{i,2}$, piecewise smooth with a kink at $x_{i,2} = x_{i,1}$. Define the frictionless target positions adjusted for the fee wedge:

$$x_{i,2}^{\pm} = \frac{\mu_{i,2} - P_2 \mp rP_2}{H\sigma_{i,2}^2}.$$

The optimum depends on whether the investor's "desired mispricing" relative to the existing position exceeds the fee-induced threshold.

Proposition 4.8 (No-trade region at date 2)**.** *Define $\Delta_2 := rP_2$. The optimal date-2 position is*

$$x_{i,2}^{\star} = \begin{cases} x_{i,2}^{+} & \text{if } \mu_{i,2} - P_2 - H\sigma_{i,2}^2 x_{i,1} > \Delta_2, \\ x_{i,1} & \text{if } |\mu_{i,2} - P_2 - H\sigma_{i,2}^2 x_{i,1}| \le \Delta_2, \\ x_{i,2}^{-} & \text{if } \mu_{i,2} - P_2 - H\sigma_{i,2}^2 x_{i,1} < -\Delta_2. \end{cases}$$

Proof. Compare the left and right derivatives of the objective at $x_{i,2} = x_{i,1}$. The left derivative equals $\mu_{i,2} - P_2 - H\sigma_{i,2}^2 x_{i,1} + rP_2$; the right derivative equals $\mu_{i,2} - P_2 - H\sigma_{i,2}^2 x_{i,1} - rP_2$. The kink is optimal if and only if the left derivative is nonnegative and the right derivative is nonpositive, which yields the inaction condition. Outside the band, the first-order condition in the smooth region gives $x_{i,2}^{\pm}$. □

The no-trade band is centered at the existing position $x_{i,1}$, not at zero: an investor who acquired a large position at date 1 requires a correspondingly larger mispricing signal to trade further, creating path dependence in portfolio dynamics (Figure 13).

At date 1, investors start with zero holdings (Assumption 4.1), which simplifies the no-trade region. The investor solves

$$\max_{x_{i,1}} \quad \mathbb{E}[V_{i,2}(x_{i,1}) \mid I_{i,1}] - rP_1|x_{i,1}|, \tag{26}$$

where $V_{i,2}$ is the value function at date 2. The continuation payoff introduces an option value: the investor's choice of $x_{i,1}$ affects the center of the date-2 inaction band. The date-1 problem nonetheless retains an inaction region centered at zero.

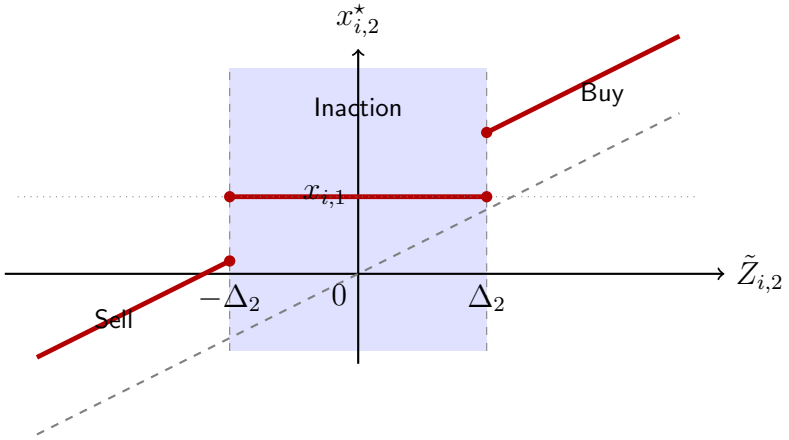

Figure 13: Optimal trading policy at date 2. The x-axis is the adjusted mispricing $\tilde{Z}_{i,2} = \mu_{i,2} - P_2 - H\sigma^2_{i,2}x_{i,1}$. The dashed gray line shows frictionless demand; the solid red line shows the optimal policy with proportional fees. In the shaded inaction region $|\tilde{Z}_{i,2}| \leq \Delta_2$, the investor retains the date-1 position $x_{i,1}$.

Proposition 4.9 (No-trade condition at date 1)**.** *Let $G_{i,1}(x) := \mathbb{E}[V_{i,2}(x) \mid I_{i,1}]$ denote the (concave) expected continuation value. Then $x^\star_{i,1} = 0$ whenever*

$$0 \in \partial G_{i,1}(0) + [-rP_1,\ rP_1].$$

If $G_{i,1}$ is differentiable at 0, this condition is equivalent to $|G'_{i,1}(0)| \leq rP_1$. The no-trade region expands in r.

Appendix E provides a formal proof using subgradient optimality conditions.

4.3 Trading Volume

Aggregate trading volume depends on the distribution of investor mispricing relative to the inaction threshold. Date-t trading volume is

$$V_t := P_t \int_0^1 |x^\star_{i,t} - x_{i,t-1}|\, di. \tag{27}$$

Dollar volume is the economically relevant quantity for founder revenue ($R_t = rV_t$). The *adjusted mispricing* for investor i at date t is

$$\tilde{Z}_{i,t} := \mu_{i,t} - P_t - H\sigma^2_{i,t}x_{i,t-1}, \tag{28}$$

which accounts for the existing position $x_{i,t-1}$ in determining trading incentives. Under the Gaussian structure and heterogeneous interpretation, $\tilde{Z}_{i,t}$ is Normally distributed across investors. At date 1, $x_{i,0} = 0$ (Assumption 4.1), so $\tilde{Z}_{i,1} = Z_{i,1}$ and the cross-sectional variance equals $\sigma^2_{Z,1}$ from

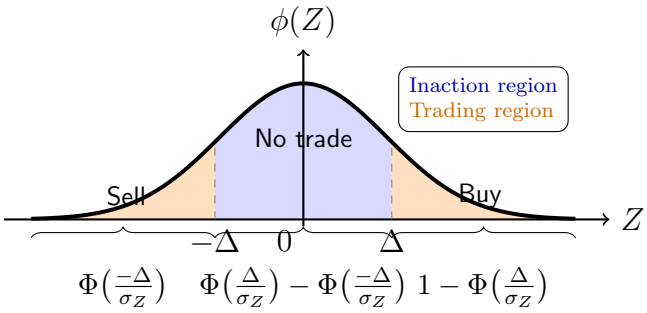

Figure 14: Distribution of adjusted mispricing $\tilde{Z} = \mu_{i,t} - P_t - H\sigma^2_{i,t}x_{i,t-1}$. Investors in the inaction region $|\tilde{Z}| \leq \Delta$ (blue) do not trade. Volume comes entirely from the tails $|\tilde{Z}| > \Delta$ (orange), where investors buy or sell to attenuated positions.

(24). At date 2, the position-adjustment term $H\sigma^2_{i,2}x_{i,1}$ varies across investors, adding covariance terms; we use $\sigma^2_{Z,2}$ as the leading-order approximation and note that the comparative statics below are evaluated under this approximation. Because investors on average hold the asset and u_t shifts aggregate demand, the distribution is *non-central*.

Proposition 4.10 (Volume representation)**.** *Using the adjusted mispricing* (28)*, date-2 expected dollar volume is*

$$V_2 = \frac{P_2}{H\sigma^2_{i,2}} \mathbb{E}\left[(|\tilde{Z}_{i,2}| - \Delta_2)_+\right], \tag{29}$$

where $(x)_+ := \max\{x, 0\}$ *and* $\Delta_2 = rP_2$ *is the inaction threshold. For* $\tilde{Z} \sim \mathcal{N}(\tilde{m}, s^2)$ *and threshold* $\Delta \geq 0$*, the truncated expectation has the closed form*

$$\mathbb{E}\left[(|\tilde{Z}| - \Delta)_+\right] = s[\phi(a) + \phi(b)] + \tilde{m}[\Phi(b) - \Phi(a)] - \Delta[2 - \Phi(a) - \Phi(b)], \tag{30}$$

where $a = (\Delta - \tilde{m})/s$*,* $b = (\Delta + \tilde{m})/s$*, and* ϕ*,* Φ *denote the standard Normal pdf and cdf.*

Formula (30) generalizes the standard result for centered Normal variables to the non-central case. Non-centrality arises because $\tilde{Z}_{i,t}$ has nonzero mean: investors on average hold the asset, and noise trader demand u_t shifts the aggregate clearing condition. Appendix E provides the proof using standard truncated Normal identities.

Figure 14 illustrates.

Section 3.4 documents that the empirical response runs through transaction counts rather than prices. The heterogeneous-interpretation model provides a natural mechanism: because investors interpret the same public signal differently, more precise creator communication *widens* the cross-sectional dispersion of posterior beliefs (equation 24), pushing more mass into the trading tails. In the standard REE, by contrast, higher public precision compresses dispersion and therefore weakens this extensive-margin trading motive. The model maps each period's creator communication to a Gaussian signal with precision κ_t. In the data, the regressor is weekly tweet count. The qualitative mapping treats more tweets as higher aggregate disclosure precision; a formal

microfoundation (e.g., each tweet as an independent signal with precision κ_0, so $\kappa_t = n_t\kappa_0$) would preserve the sign predictions but is not pursued here.

Evaluating unconditionally over aggregate states (so $\mathbb{E}[\tilde{Z}_{i,t}] = 0$ and the cross-sectional variance is $\sigma^2_{Z,t}$), the expected transaction count is the fraction of investors who trade:

$$N_t = \Pr(|\tilde{Z}_{i,t}| > \Delta_t) = 2\left[1 - \Phi\left(\frac{\Delta_t}{\sigma_{Z,t}}\right)\right].$$

Since N_t is increasing in $\sigma_{Z,t}$ (in the data, transaction counts vary with the pool of potential traders per collection; collection fixed effects absorb permanent differences in holder-base size) and Appendix D shows that $\sigma_{Z,t}$ is increasing in current disclosure once current effective precision exceeds a critical threshold determined by inherited disagreement and cumulative precision, we obtain that expected volume is increasing in public precision under disagreement. In the standard REE, the opposite dispersion channel applies: higher public precision compresses disagreement and weakens the extensive-margin response.

The volume response operates primarily through the *transaction-count (extensive) margin*. In the model, disclosure directly affects trading by changing the fraction of investors whose perceived mispricing crosses the fee-induced threshold, while expected prices move only through the posterior-variance term $\mathbb{E}[P_t] = \theta - H\sigma^2_{i,t}$. The mechanism is therefore transaction-dominant rather than price-dominant. The empirical finding of approximately two-thirds (Table 11) is a less extreme counterpart: the model predicts the response is entirely extensive-margin, while the data show a positive but smaller intensive-margin (price) component. The nonzero price elasticity (0.090) likely reflects channels outside the model—such as price impact from order flow, liquidity effects, or attention-driven revaluation—that the CARA–Normal framework abstracts from.

The cross-partial derivative of expected volume with respect to public precision and the fee rate is positive ($\partial^2\mathbb{E}[V_t]/(\partial\kappa_t\,\partial r) > 0$) in the empirically relevant parameter range. Higher fees *steepen* the volume–disclosure relationship: the inaction threshold places more investors "on the margin," so disagreement-induced widening of the Z-distribution has a larger effect. Conversely, when fees decrease, the volume–disclosure link attenuates; in the limit $r \to 0$, all investors trade regardless of $\sigma_{Z,t}$, and $\partial V_t/\partial\kappa_t \to 0$. The empirical test (Table 7) interacts tweet count with the contractual royalty rate and enforcement-regime indicators, producing a triple interaction. This is a second-order implication of the cross-partial: if $\partial^2 V/(\partial\kappa\,\partial r) > 0$, then reducing effective r (weaker enforcement) should flatten the fee-heterogeneity gradient, which is the pattern observed.

The model's limiting prediction ($\partial V_t/\partial\kappa_t \to 0$ as $r \to 0$) is not literally tested, since effective per-trade costs never reached zero: Ethereum gas fees and residual slippage persisted throughout the post-erosion period. The positive but attenuated post-erosion semi-elasticity (0.56%, Table 5) is consistent with a reduced but nonzero effective wedge.

These comparative statics match the data: higher congestion and stricter enforcement raise the

effective wedge and amplify the information-to-volume link; optional royalties shrink the wedge and flatten it; Solana provides a low-fee benchmark where the gas component is negligible.

Differentiating with respect to the fee rate yields an unambiguous sign:

$$\frac{\partial V_2}{\partial r} = -\frac{P_2^2}{H\sigma_{i,2}^2} \Pr(|\tilde{Z}_{i,2}| > \Delta_2) < 0. \tag{31}$$

Higher fees reduce volume both directly (wider inaction bands) and indirectly (fewer investors clear the threshold). The sign holds for all parameter values.

Private precision has an ambiguous effect on volume. Define

$$\ell(\tau_2) := \tau_2 \sigma_{i,2}^2, \tag{32}$$

which captures how private precision scales belief dispersion. In the standard CARA–Normal REE, *higher private precision* has an ambiguous effect on volume: it widens the cross-section of posterior means through investor-specific private signals (raising trading motives) but also reduces posterior variance (shrinking position sizes). The net effect depends on how these forces balance. In the disagreement model studied here, by contrast, sufficiently informative creator communication raises expected volume through the heterogeneous-interpretation term even while posterior variance falls. Fixed-fee alternatives therefore eliminate deadweight loss at revenue equivalence.

The model therefore fits the empirical findings through a single sequence of comparative statics. More informative creator communication raises cross-sectional dispersion in posterior beliefs, which increases the fraction of investors whose perceived mispricing exceeds the fee-induced inaction band; this generates a positive tweet–volume relation and makes the response operate mainly through transaction counts rather than prices. Because the trading threshold is proportional to the wedge $\Delta_t = rP_t$, higher royalties and higher congestion make more investors marginal and therefore steepen the communication–volume gradient, while weaker royalty enforcement lowers the wedge and flattens that gradient. The same structure then maps into the policy results: royalties reward volume and thus tilt disclosure toward periods where disagreement can stimulate trade, whereas the benchmark upfront membership counterfactual removes the trading wedge, induces front-loaded disclosure, and raises investor welfare.

4.4 Optimal Disclosure Timing

The founder allocates disclosure precision across periods to maximize royalty revenue. The problem follows the spirit of Bayesian persuasion (Kamenica and Gentzkow, 2011), though the receiver's action—trading—is distorted by proportional fees and filtered through heterogeneous interpretation.

The founder solves

$$\max_{\kappa_1,\kappa_2\geq 0} \Pi_F(\kappa_1,\kappa_2) := r(V_1+V_2) \quad \text{subject to} \quad \kappa_1+\kappa_2=\bar{\kappa}. \tag{33}$$

Total royalty revenue equals the fee rate times aggregate volume across both periods.

Section 4.3 gives

$$V_t = \frac{P_t}{H\sigma_{i,t}^2}\,\mathbb{E}\left[(|\tilde{Z}_{i,t}|-\Delta_t)_+\right], \qquad \Delta_t = rP_t, \tag{34}$$

where $\tilde{Z}_{i,t}$ is the adjusted mispricing defined in (28). At date 2 this representation follows from the piecewise-linear optimum in Proposition 4.8. At date 1, the investor solves a dynamic problem with continuation value (equation 26), so (34) is a myopic approximation that abstracts from the option value of future rebalancing; the option value widens the true date-1 inaction band, making date-1 volume weakly lower than the myopic formula suggests. Disclosure affects volume through both the posterior variance and the cross-sectional dispersion of beliefs. From (14),

$$\frac{\partial\sigma_{i,t}^2}{\partial\kappa_t} = -(\sigma_{i,t}^2)^2\cdot\frac{\psi_t^2}{(\kappa_t+\psi_t)^2} < 0. \tag{35}$$

More disclosure reduces residual uncertainty. Under disagreement, higher κ_t can also widen the cross-sectional dispersion $\sigma_{Z,t}$; Appendix D gives the exact threshold condition under which this disagreement effect dominates the compression from higher posterior precision.

Differentiating (34) with respect to κ_t and applying the chain rule:

$$\frac{\partial V_t}{\partial\kappa_t} = \frac{P_t}{H\sigma_{i,t}^2}\left[\frac{\partial}{\partial\kappa_t}\mathbb{E}\left[(|\tilde{Z}_{i,t}|-\Delta_t)_+\right] - \frac{\partial\ln(H\sigma_{i,t}^2)}{\partial\kappa_t}\cdot\mathbb{E}\left[(|\tilde{Z}_{i,t}|-\Delta_t)_+\right]\right]. \tag{36}$$

The first term captures how disclosure shifts the mispricing distribution through disagreement; the second captures the scale effect through posterior variance. Under heterogeneous interpretation, the first term dominates once current effective precision is above the threshold characterized in Appendix D, yielding $\partial V_t/\partial\kappa_t > 0$.

Define period-t royalty revenue as

$$R_t := rV_t = rP_t\,\mathbb{E}\big[|x_{i,t}^\star - x_{i,t-1}|\big], \tag{37}$$

so that $\Pi_F = R_1 + R_2$. Interior solutions satisfy the first-order condition

$$\frac{\partial R_1}{\partial\kappa_1} = \frac{\partial R_2}{\partial\kappa_2}. \tag{38}$$

The founder equates the marginal revenue contribution of disclosure across periods.

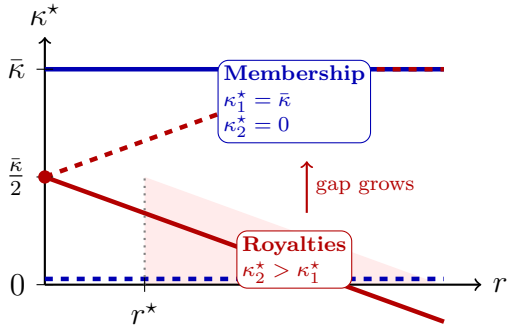

Figure 15: Stylized disclosure timing under alternative fee structures. Under royalties (red), higher r can tilt the optimal split toward back-loading ($\kappa_2^\star > \kappa_1^\star$). Under the benchmark upfront membership fee (blue), full front-loading ($\kappa_1^\star = \bar{\kappa}$, $\kappa_2^\star = 0$) is optimal.

Define the relative revenue ratio and the elasticity ratio:

$$\Psi(r, \sigma_{u,2}^2) := \frac{P_2}{P_1} \cdot \frac{\mathbb{E}[|x_{i,2}^\star - x_{i,1}|]}{\mathbb{E}[|x_{i,1}^\star|]}, \tag{39}$$

$$\Theta(r, \sigma_{u,2}^2) := \frac{\omega_1(r)}{\omega_2(r)}, \quad \text{where } \omega_t := \frac{\partial \ln V_t}{\partial \ln \kappa_t}. \tag{40}$$

The ratio Ψ measures the relative size of date-2 versus date-1 trading (in expectation); Θ measures the relative responsiveness of volume to disclosure. When Ψ is high and/or Θ is low, the founder has incentives to allocate more disclosure to date 2 (back-loading).

Remark 4.11 (Back-loading under royalties)**.** *Under royalties, the founder's revenue is proportional to volume, which is increasing in disagreement-induced dispersion. This creates an incentive to concentrate disclosure in later periods where the marginal volume gain is largest—a back-loading motive. However, the concavity of κ_t^{eff} in κ_t (Remark 4.2) creates a countervailing precision-loss effect, so the optimal disclosure split depends on parameters. The back-loading tendency is established numerically at the baseline calibration; a formal analytical characterization is left for future work. Figure 15 illustrates how the disclosure split tilts toward the later period under royalties, while Proposition 4.13 shows that under the benchmark upfront ad-valorem membership contract, full front-loading is optimal.*

The back-loading prediction is testable: under royalties, creators should release more information later in a project's lifecycle. The present data set records tweet timestamps and could in principle test whether disclosure timing shifted when fee enforcement weakened, though we do not pursue this test here.

4.5 Welfare Analysis

We study welfare improvements from changing the fee structure. Comparing proportional royalties with revenue-equivalent membership contracts reveals two mechanisms: eliminating trading-margin distortions and changing creators' disclosure incentives. In the baseline CARA framework with heterogeneous interpretation, the welfare gain comes from eliminating the per-trade wedge (Proposition 4.14). Fee structure also governs whether creators disclose information early (to resolve uncertainty and raise prices) or late (to stimulate trading volume and generate disagreement). Under the ERU preference extension in the Online Appendix, the benchmark shift from royalties to the upfront membership contract provides an additional welfare gain by aligning information release with efficient uncertainty resolution.

Royalties. Investor welfare net of fee transfers is

$$W(r) = \mathbb{E}[f] - \frac{H}{2}\mathbb{E}\left[\sigma_{i,2}^2 (x_{i,2}^\star)^2\right] - rP_2\mathbb{E}\left[|x_{i,2}^\star - x_{i,1}^\star|\right] - rP_1\mathbb{E}\left[|x_{i,1}^\star|\right]. \tag{41}$$

The first term is the expected asset payoff; the second is the risk-bearing cost (investors hold risky positions); the third and fourth are expected fee payments. Expected founder revenue under royalties is $\mathbb{E}[r(V_1 + V_2)]$.

Consider instead a revenue-equivalent membership fee M satisfying (44), collected once rather than per trade. The per-trade wedge vanishes ($\Delta_t = 0$), and investors trade to their frictionless optima.

Deadweight loss from proportional fees. Consider a generic period with wedge $\Delta = rP$ and frictionless mispricing Z. Define

$$x^0 = \frac{Z}{H\sigma^2}, \qquad x^\star = \frac{Z - \operatorname{sgn}(Z)\Delta}{H\sigma^2} \cdot \mathbf{1}\{|Z| > \Delta\}.$$

The frictionless position x^0 fully exploits the information in Z; the frictional position $x^\star$ either equals the entry position (if $|Z| \le \Delta$) or adjusts partially (if $|Z| > \Delta$).

Proposition 4.12 (Welfare loss from proportional fees)**.** *The per-period welfare loss from the trading wedge is*

$$\Delta W = \frac{1}{2H\sigma^2}\mathbb{E}\left[Z^2\mathbf{1}\{|Z| \le \Delta\} + (2|Z|\Delta - \Delta^2)\mathbf{1}\{|Z| > \Delta\}\right] > 0. \tag{42}$$

Proof. Fix Z and consider the concave objective

$$U_\Delta(Z) := \max_{x \in \mathbb{R}} \left\{ Zx - \frac{H\sigma^2}{2}x^2 - \Delta|x| \right\}.$$

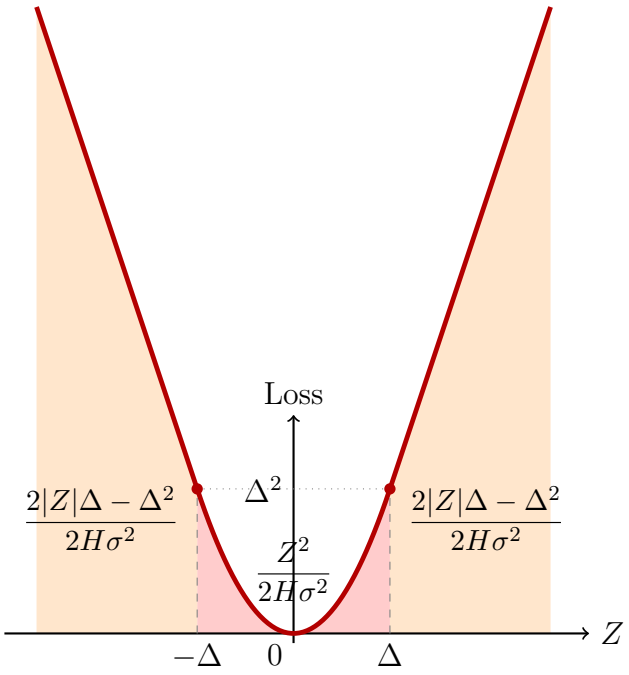

Figure 16: Welfare loss as a function of mispricing Z. In the inaction region $|Z| \leq \Delta$ (red), the entire surplus from the frictionless position is lost. In the trading region $|Z| > \Delta$ (orange), partial surplus is lost due to attenuated positions.

Without the wedge ($\Delta = 0$), the optimizer is $x^0 = Z/(H\sigma^2)$ and $U_0(Z) = Z^2/(2H\sigma^2)$. With $\Delta > 0$, the maximizer is $x^\star = 0$ if $|Z| \leq \Delta$ and $x^\star = (Z - \mathrm{sgn}(Z)\Delta)/(H\sigma^2)$ if $|Z| > \Delta$. Evaluating the objective at $x^\star$ gives

$$U_\Delta(Z) = \begin{cases} 0, & |Z| \leq \Delta, \\ \dfrac{(|Z| - \Delta)^2}{2H\sigma^2}, & |Z| > \Delta. \end{cases}$$

Therefore the state-by-state welfare loss is

$$U_0(Z) - U_\Delta(Z) = \frac{1}{2H\sigma^2}\Big[Z^2 \mathbf{1}\{|Z| \leq \Delta\} + \big(2|Z|\Delta - \Delta^2\big)\mathbf{1}\{|Z| > \Delta\}\Big].$$

Taking expectations yields (42). Strict positivity follows for $\Delta > 0$ since the integrand is almost surely nonnegative and strictly positive on a set of positive probability for any continuously distributed Z. □

Figure 16 illustrates.

Participation margins. The comparison holds the set of market participants fixed. In practice, a membership fee set too high could deter entry, eroding the welfare gains from eliminating per-trade distortions. The analysis provides an upper bound on the welfare gain, conditional on full participation under both regimes.

4.6 Membership Fee Counterfactual

Consider a family of ad-valorem membership fees proportional to an average of equilibrium prices. This structure preserves founder revenue while removing the inaction bands that suppress trading. The main policy experiment below focuses on the *upfront* benchmark in which the fee is paid at date 1 and indexed to the date-1 market price.

Ad-valorem membership fees. The founder charges a deterministic membership fee, indexed to expected market conditions via

$$M = m \cdot \mathbb{E}[\bar{P}], \qquad \bar{P} := w_1 P_1 + w_2 P_2, \tag{43}$$

where $w_1, w_2 \geq 0$ with $w_1 + w_2 = 1$. The fee is proportional to an expected weighted average of prices, indexing the founder's claim to market conditions while remaining deterministic—eliminating fee-payment risk for CARA investors.

Revenue equivalence requires

$$M = m \cdot \mathbb{E}[\bar{P}] = \mathbb{E}[r(V_1 + V_2)]. \tag{44}$$

This pins down m as a function of r and the royalty-regime revenue. The benchmark upfront contract sets $(w_1, w_2) = (1, 0)$, so $M = m\,\mathbb{E}[P_1]$.

Ad-valorem membership eliminates the per-trade wedge: investors face $\Delta_t = 0$ and trade to their frictionless optima.

Optimal disclosure under membership.

Proposition 4.13 (Disclosure under upfront ad-valorem membership)**.** *Under the benchmark upfront ad-valorem membership contract $(w_1, w_2) = (1, 0)$, the optimal disclosure policy is*

$$\kappa_1^\star = \bar{\kappa}, \qquad \kappa_2^\star = 0.$$

Revenue under membership is

$$M = m \cdot \mathbb{E}[\bar{P}] = m \cdot [w_1 \mathbb{E}[P_1] + w_2 \mathbb{E}[P_2]]. \tag{45}$$

From Proposition 4.4, $\mathbb{E}[P_t] = \theta - H\sigma_{i,t}^2$. Differentiating with respect to κ_1 along the budget

constraint:

$$\frac{\partial \mathbb{E}[P_1]}{\partial \kappa_1} = H(\sigma_{i,1}^2)^2 \cdot \frac{\psi_1^2}{(\kappa_1 + \psi_1)^2} > 0, \tag{46}$$

$$\frac{d\mathbb{E}[P_2]}{d\kappa_1} = H(\sigma_{i,2}^2)^2 \left[\frac{\psi_1^2}{(\kappa_1 + \psi_1)^2} - \frac{\psi_2^2}{(\kappa_2 + \psi_2)^2} \right]. \tag{47}$$

Consequently,

$$\frac{d\mathbb{E}[\bar{P}]}{d\kappa_1} = w_1 \frac{\partial \mathbb{E}[P_1]}{\partial \kappa_1} + w_2 \frac{d\mathbb{E}[P_2]}{d\kappa_1}. \tag{48}$$

For general weighted contracts, the second term can be positive, zero, or negative, so interior disclosure splits are possible. The sharp corner result below therefore applies to the benchmark upfront contract that mirrors a one-time membership purchase at entry.

Proof of Proposition 4.13. Under the benchmark upfront contract, the founder solves $\max_{\kappa_1 \in [0,\bar{\kappa}]} \mathbb{E}[P_1]$ subject to $\kappa_2 = \bar{\kappa} - \kappa_1$. Equation (46) shows that $\partial \mathbb{E}[P_1]/\partial \kappa_1 > 0$ for every $\kappa_1 \in [0, \bar{\kappa}]$. Therefore the objective is strictly increasing in κ_1, so the unique optimum is the corner solution $\kappa_1^\star = \bar{\kappa}$ and $\kappa_2^\star = 0$. □

Under the benchmark upfront membership contract, the founder strictly prefers to front-load disclosure. Early disclosure raises P_1 by reducing uncertainty, so the founder reveals all available information immediately. For more general weighted contracts, equation (48) shows how the incentive depends on the relative weight on date-1 versus date-2 prices.

The sharp front-loading result depends on the benchmark upfront weighting $(w_1, w_2) = (1, 0)$. For the equal-weight contract $(w_1, w_2) = (1/2, 1/2)$, equation (48) can produce interior solutions. For the end-loaded contract $(w_1, w_2) = (0, 1)$, the founder maximizes $\mathbb{E}[P_2]$ alone, and by the same concavity logic, full *back*-loading ($\kappa_1 = 0$, $\kappa_2 = \bar{\kappa}$) becomes optimal—exactly reversing the disclosure prediction. The welfare result (Proposition 4.14) is robust to any weighting because it depends only on eliminating the per-trade wedge, but the disclosure-timing channel and its welfare implications under ERU preferences are specific to the upfront benchmark.

Under royalties the founder may back-load disclosure (Remark 4.11) because volume—not prices—determines revenue, and under disagreement, more disclosure generates more volume through widening belief dispersion. Fee structure thus governs disclosure incentives: royalties incentivize late disclosure to stimulate trading through disagreement, while the benchmark upfront membership fee incentivizes early disclosure to resolve uncertainty. This shift in incentives changes disclosure behavior, which in turn improves welfare by aligning information release with efficient uncertainty resolution rather than with trading-volume maximization.

Proposition 4.14 (Ad-valorem membership yields higher investor welfare)**.** *Under revenue equivalence (44) and full participation, ad-valorem membership yields higher investor welfare than proportional royalties: $W(M) > W(r)$.*

Proof. Fix a period and let Z denote the investor's frictionless mispricing and $\Delta := rP$ the proportional wedge. As in Proposition 4.12, define the net trading surplus

$$U_\Delta(Z) := \max_{x\in\mathbb{R}} \left\{ Zx - \frac{H\sigma^2}{2}x^2 - \Delta|x| \right\},$$

with optimizer $x^\star$ and note that the royalty payment equals $\Delta|x^\star|$. Under membership, the per-trade wedge is zero, so the investor attains $U_0(Z)$ and pays the fixed fee M once. At revenue equivalence, expected payments are identical across regimes, so the welfare comparison reduces to the net-of-transfer loss

$$U_0(Z) - U_\Delta(Z) - \Delta|x^\star|.$$

By the same casework as in the proof of Proposition 4.12, $x^\star = 0$ when $|Z| \le \Delta$ and $x^\star = \frac{Z - \operatorname{sgn}(Z)\Delta}{H\sigma^2}$ when $|Z| > \Delta$. Therefore,

$$U_0(Z) - U_\Delta(Z) - \Delta|x^\star| = \frac{1}{2H\sigma^2}\Big[Z^2\mathbf{1}\{|Z| \le \Delta\} + \Delta^2\mathbf{1}\{|Z| > \Delta\}\Big] > 0$$

whenever Z has a continuous distribution with $\Pr(|Z| > 0) > 0$. Taking expectations and summing across $t \in \{1, 2\}$ yields $W(M) > W(r)$. □

Remark 4.15 (Revenue equivalence and disclosure endogeneity)**.** *Revenue equivalence* (44) *equates expected payments across regimes, each evaluated at the respective optimal disclosure policy:* (κ_1^R, κ_2^R) *under royalties and* $(\bar\kappa, 0)$ *under upfront membership. The fee rate* m *therefore implicitly depends on both disclosure regimes. The welfare comparison remains valid because the state-by-state inequality* $U_0(Z) - U_\Delta(Z) - \Delta|x^\star| > 0$ *holds for every realization of* Z*, regardless of its distribution. Because the membership fee* $M = m\,\mathbb{E}[P_1]$ *is deterministic, CARA investors bear no fee-payment risk under membership, while royalty payments* $rP_t|x^\star_{i,t} - x_{i,t-1}|$ *are stochastic. This asymmetry favors membership, so the welfare ranking is reinforced by risk considerations.*

Nesting: standard REE as a special case. As interpretation precision $\psi_t \to \infty$, the heterogeneous-interpretation model collapses smoothly to the common-signal CARA–Normal noisy REE without additional price learning (i.e., maintaining $\phi_{P,t} = 0$):

(i) $\kappa_t^{\text{eff}} = \kappa_t\psi_t/(\kappa_t + \psi_t) \to \kappa_t$.

(ii) $\sigma_{i,t}^{-2} \to \sigma_f^{-2} + \sum_{s=1}^t(\tau_s + \kappa_s)$ (the standard cumulative posterior precision).

(iii) The disagreement component vanishes: $(\kappa_t^{\text{eff}})^2/\psi_t \to 0$.

(iv) $\sigma_{Z,t}^2 \to \left(\sum_{s=1}^t \tau_s\right)/(\sigma_{i,t}^{-2})^2$ (the standard formula, decreasing in cumulative public precision).

(v) The disagreement-driven public-precision effect disappears: higher public precision compresses dispersion instead of widening it, so the standard REE no longer delivers the positive extensive-margin mechanism emphasized in the paper.

(vi) The $\phi_{P,t} = 0$ restriction is maintained in the limit. In the full Admati (1985) noisy REE, investors additionally learn from equilibrium prices ($\phi_{P,t} > 0$). Restoring price learning would add a common-information channel that compresses dispersion, partially offsetting the disagreement mechanism for finite ψ_t.

This nesting shows that the disagreement model generalizes the standard framework through a single parameter (ψ_t), with the common-signal model (with $\phi_P = 0$) recovered in the limit.

CARA preferences are neutral to the timing of uncertainty resolution. An Online Appendix (in preparation for the disagreement-model specification) extends the framework to preferences for early resolution of uncertainty (ERU) by augmenting date-1 risk aversion. ERU widens the date-1 inaction region and amplifies the divergence between fee regimes: royalties push disclosure later, while the benchmark upfront membership contract pushes disclosure earlier. Membership yields higher investor welfare than royalties under both preference specifications.

5 Conclusion

Creator tweets predict NFT trading volume through the extensive margin, and the link weakens when proportional royalty enforcement declines. A three-regime design on Ethereum and cross-chain replication on Solana provide convergent evidence consistent with a fee-structure channel above and beyond aggregate shocks. A noisy rational expectations model with heterogeneous interpretation of public signals shows that proportional fees create inaction regions and that disagreement—arising because different investors interpret the same creator communication differently—widens belief dispersion and amplifies the information-volume gradient through the transaction-count margin. The model nests the standard common-signal REE as a special case. The central policy insight is that revenue-equivalent membership contracts improve welfare by eliminating trading distortions, and that the benchmark upfront membership counterfactual additionally shifts disclosure incentives from late disclosure (to stimulate trading) to early disclosure (to raise prices by resolving uncertainty). These welfare comparisons are conditional on full participation and revenue equivalence.

The results bear on platform design in digital asset markets. Per-trade royalties distort marginal trading decisions by widening inaction regions and also distort disclosure incentives. NFT markets lack disclosure regulation, so creators time communication to maximize fee revenue. Fee structure determines the resulting behavior: proportional royalties incentivize back-loaded disclosure that stimulates trades; the benchmark upfront membership contract incentivizes front-loaded

disclosure that resolves uncertainty. In the baseline CARA model with heterogeneous interpretation, welfare improvements come from eliminating the per-trade wedge; under ERU preferences, the disclosure-timing shift provides an additional welfare gain. The mechanism extends beyond NFTs to any setting where intermediary income is tied to transaction flow and investors interpret public disclosures heterogeneously.

This also clarifies why the royalty-enforcement episodes are informative policy experiments rather than just institutional anecdotes. In the model, changing fee enforcement changes the slope linking creator communication to trading, not merely the level of trading costs. The observed attenuation of the tweet–volume semi-elasticity when enforcement weakens is therefore evidence about market design's effect on information transmission. That is what allows the paper to use the policy variation to speak to counterfactual platform design: moving from royalties to membership fees changes not only who pays and when, but also how much creators disclose, when they disclose it, and how much welfare is lost to wedge-induced non-trading.

Identification is limited by policy event endogeneity and the FTX timing coincidence. Blur's airdrop program incentivized wash trading that may independently attenuate the semi-elasticity. Within-week endogeneity means the semi-elasticity captures an association rather than a clean causal effect, even with hourly evidence. The welfare comparison assumes CARA-Gaussian preferences, full participation, and revenue equivalence between fee structures. The disclosure-timing predictions (back-loading under royalties, front-loading under membership) are derived theoretically but not tested empirically. The data contain tweet timestamps that could in principle test whether creator disclosure behavior changed when fee enforcement eroded; this remains for future work.

References

Admati, Anat R., "A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets," *Econometrica*, 1985, *53* (3), 629–657.

Amihud, Yakov, "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects," *Journal of Financial Markets*, 2002, *5* (1), 31–56.

Antweiler, Werner and Murray Z. Frank, "Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards," *Journal of Finance*, 2004, *59* (3), 1259–1294.

Araci, Dogu, "FinBERT: Financial Sentiment Analysis with Pre-Trained Language Models," *arXiv preprint arXiv:1908.10063*, 2019.

Armstrong, Mark, "Competition in Two-Sided Markets," *RAND Journal of Economics*, 2006, *37* (3), 668–691.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Blankespoor, Elizabeth, Gregory S. Miller, and Hal D. White, "The Role of Dissemination in Market Liquidity: Evidence from Firms' Use of Twitter," *The Accounting Review*, 2014, *89* (1), 79–112.

Budish, Eric, "Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains," *Quarterly Journal of Economics*, 2025, *140* (1), 1–58.

Constantinides, George M., "Capital Market Equilibrium with Transaction Costs," *Journal of Political Economy*, 1986, *94* (4), 842–862.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, "In Search of Attention," *Journal of Finance*, 2011, *66* (5), 1461–1499.

Easley, David, Maureen O'Hara, and Soumya Basu, "From Mining to Markets: The Evolution of Bitcoin Transaction Fees," *Journal of Financial Economics*, 2019, *134* (1), 91–109.

Grossman, Sanford J. and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, *70* (3), 393–408.

Harris, Milton and Artur Raviv, "Differences of Opinion Make a Horse Race," *Review of Financial Studies*, 1993, *6* (3), 473–506.

Huberman, Gur, Jacob D. Leshno, and Ciamac C. Moallemi, "Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System," *Review of Economic Studies*, 2021, *88* (6), 3011–3062.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kandel, Eugene and Neil D. Pearson, "Differential Interpretation of Public Signals and Trade in Speculative Markets," *Journal of Political Economy*, 1995, *103* (4), 831–872.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315–1335.

Loughran, Tim and Bill McDonald, "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *Journal of Finance*, 2011, *66* (1), 35–65.

Nadini, Matthieu, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli, "Mapping the NFT Revolution: Market Trends, Trade Networks, and Visual Features," *Scientific Reports*, 2021, *11*, 20902.

Petersen, Mitchell A., "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches," *Review of Financial Studies*, 2009, *22* (1), 435–480.

Reimers, Nils and Iryna Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in "Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing" 2019, pp. 3982–3992.

Rochet, Jean-Charles and Jean Tirole, "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 2006, *37* (3), 645–667.

Rysman, Marc, "The Economics of Two-Sided Markets," *Journal of Economic Perspectives*, 2009, *23* (3), 125–143.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Tetlock, Paul C., "Giving Content to Investor Sentiment: The Role of Media in the Stock Market," *Journal of Finance*, 2007, *62* (3), 1139–1168.

Vayanos, Dimitri and Jiang Wang, "Market Liquidity—Theory and Empirical Evidence," in "Handbook of the Economics of Finance," Vol. 2, Elsevier, 2012, pp. 1289–1361.

Vives, Xavier, *Information and Learning in Markets: The Impact of Market Microstructure*, Princeton, NJ: Princeton University Press, 2008.

An Online Appendix (in preparation) will contain additional empirical results, time-series diagnostics, hourly temporal precedence evidence, an early-resolution-of-uncertainty (ERU) preference extension, calibration exercises mapping the theory to both blockchain panels, and numerical verification checks.

A Sample Construction

A.1 Data Sources

Blockchain transaction data come from public APIs (Etherscan for Ethereum, Solscan for Solana) and yield weekly trading outcomes—USD volume on both chains, plus transaction counts and average prices on Ethereum—along with contractual royalty rates.

Twitter data come from the Twitter API v2. We identify official collection accounts through (i) links in collection metadata on OpenSea/Magic Eden, (ii) verified accounts matching collection names, and (iii) manual cross-referencing against Discord announcements.

A.2 Sample Selection

We apply the following filters:

Ethereum:

1. Collections with at least 1,000 lifetime secondary transactions (ensures sufficient trading activity)
2. Collections with identifiable, active Twitter accounts (at least 10 tweets during sample period)
3. Collections launched before January 2023 (ensures sufficient pre/post Blur observations)
4. Excluding derivative/wrapped collections and fractionalized NFTs

These filters yield 169 Ethereum collections.

Solana:

1. Collections with at least 500 lifetime secondary transactions
2. Collections with identifiable, active Twitter accounts
3. Collections launched before October 2022

These filters yield 83 Solana collections.

A.3 Variable Construction

Trading variables are aggregated to weekly frequency (Tuesday–Monday). Volume is converted to USD using weekly average ETH/USD or SOL/USD rates from Yahoo Finance to avoid endogeneity with crypto market conditions.

Tweet counts include original tweets and quote tweets from official accounts, excluding retweets and replies.

Tweet informativeness. We construct a composite informativeness score at the individual tweet level from five text-based dimensions:

(i) *Sentiment magnitude*: the absolute value of financial sentiment, computed using the FinBERT language model (Araci, 2019), which is pre-trained on financial text and distinguishes positive, negative, and neutral sentiment. We use the absolute value because strong sentiment in either direction signals that the tweet conveys a definite assessment rather than vague commentary.

(ii) *Specificity*: the density of named entities (proper nouns, project names, partner organizations) and numeric references (dates, quantities, prices) in the tweet text, identified using spaCy's named-entity recognition pipeline. Specific, factual tweets score higher than generic promotional language.

(iii) *Novelty*: the semantic distance between the current tweet and the collection's recent communication, measured as one minus the maximum cosine similarity between Sentence-BERT embeddings (Reimers and Gurevych, 2019) of the current tweet and all tweets posted by the same collection in the preceding 30 days. This component captures whether a tweet introduces genuinely new information or repeats prior messaging.

(iv) *Certainty*: one minus the proportion of Loughran and McDonald (2011) uncertainty words (*may*, *might*, *possibly*, *approximately*, etc.) in the tweet text. Tweets that communicate with conviction score higher.

(v) *Announcement classification*: a binary indicator equal to one if a zero-shot classification model assigns the tweet to the "announcement" category (over alternatives including "promotion," "engagement," and "community"). This component identifies concrete project updates such as roadmap milestones, partnership disclosures, and feature launches.

The composite score is a weighted average of the five components:

$$\text{Informativeness}_j = 0.20\, S_j + 0.25\, P_j + 0.20\, N_j + 0.15\, C_j + 0.20\, \mathbf{1}\{A_j\},$$

where S is sentiment magnitude, P specificity, N novelty, C certainty, and A announcement classification. Specificity receives the largest weight because factual content is most directly informative

about value. The collection-week measure averages across all tweets in week t:

$$\overline{\text{Info}}_{it} = \frac{1}{|\mathcal{T}_{it}|} \sum_{j \in \mathcal{T}_{it}} \text{Informativeness}_j,$$

For collection-weeks with no tweets, the variable is undefined; informativeness regressions restrict to $|\mathcal{T}_{it}| > 0$. The replication package includes a keyword-frequency alternative for environments without the NLP libraries; all reported results use the full transformer-based pipeline.

B Additional Empirical Results

This appendix collects supplementary figures and tables referenced in the main text. Figures 17–20 display variable distributions and time series for both blockchains. Figures 21 and 22 report event-study estimates; Figure 23 plots informativeness partial correlations. The remaining tables present Solana-specific analyses paralleling the Ethereum results, placebo-break diagnostics, and a binned residual scatter.

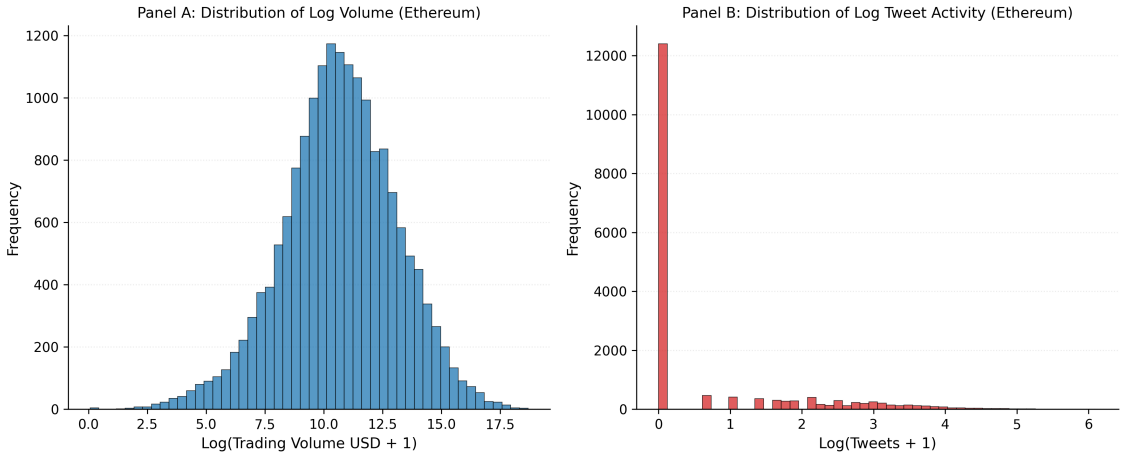

Figure 17: Distribution of Trading Volume and Tweet Activity (Ethereum)

Notes: Panel A displays the distribution of $\log(1 + \text{Volume})$ and Panel B shows the distribution of $\log(1 + \text{Tweets})$. Log transformations reduce right-skew and outlier sensitivity, motivating the log-linear specifications used throughout.

C Robustness Checks

Tables 25 and 26 report robustness checks for both blockchains. Panel B.1 excludes the FTX period (November 2022–January 2023). Panel B.2 winsorizes extreme weekly volume observations. Panel B.3 estimates the baseline in first differences to mitigate spurious-regression concerns. Panel B.4 reports HC1 standard errors to illustrate the role of within-collection dependence.

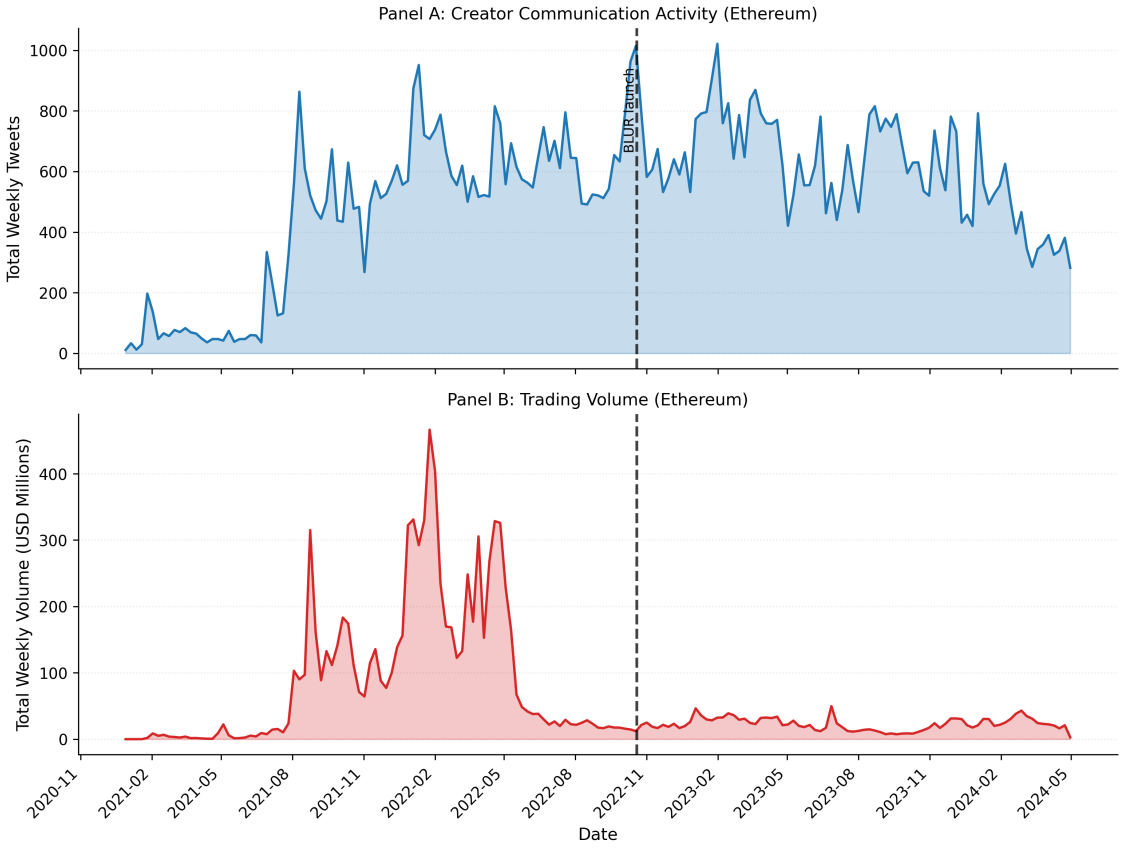

Figure 18: Time Series of Creator Activity and Trading Volume (Ethereum)

Notes: Panel A shows aggregate weekly tweet activity and Panel B shows aggregate weekly trading volume across all Ethereum NFT collections in the sample. The vertical line marks Blur's launch (October 19, 2022). Both series exhibit seasonal variation and co-movement with broader crypto market conditions.

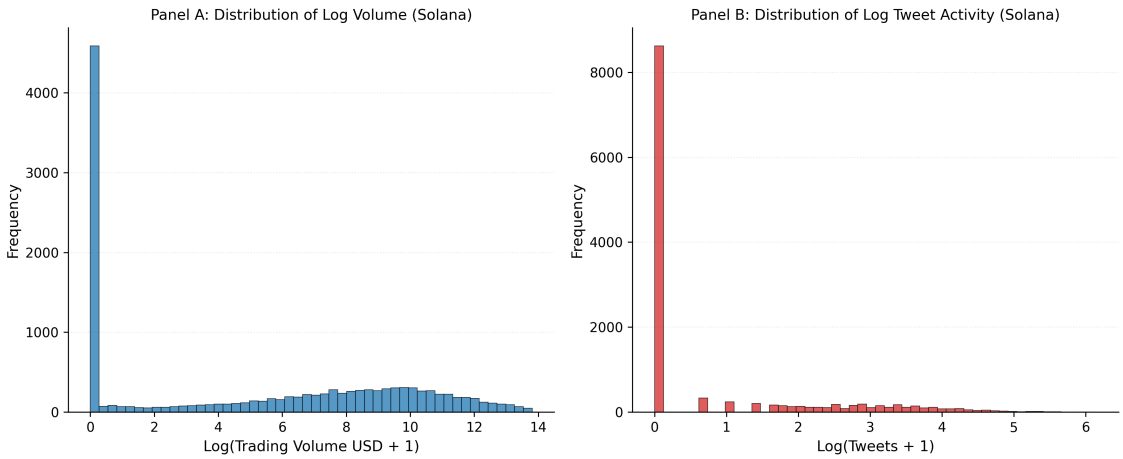

Figure 19: Distribution of Trading Volume and Tweet Activity (Solana)

Notes: Panel A displays the distribution of $\log(1 + \text{Volume})$ and Panel B shows the distribution of $\log(1 + \text{Tweets})$ for Solana collections. Log transformations reduce right-skew and outlier sensitivity, motivating the log-linear specifications used throughout.

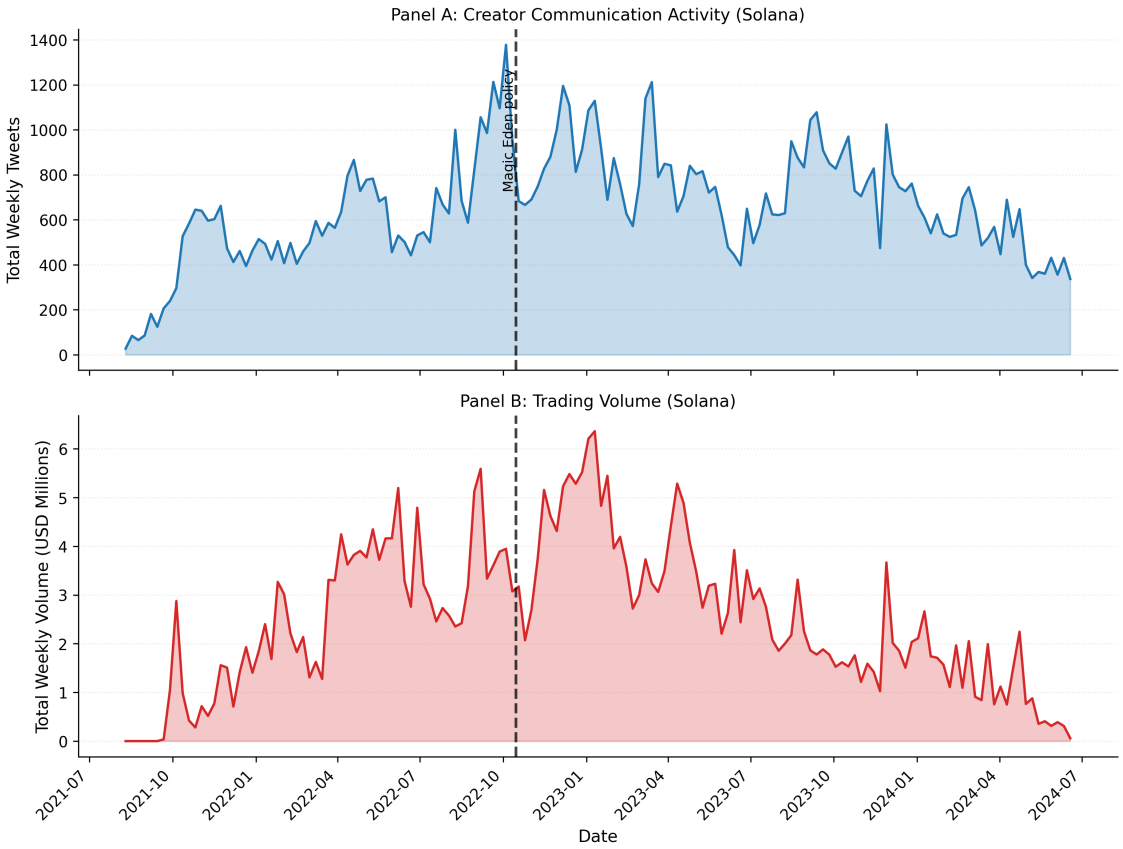

Figure 20: Time Series of Creator Activity and Trading Volume (Solana)

Notes: Panel A shows aggregate weekly tweet activity and Panel B shows aggregate weekly trading volume across all Solana NFT collections in the sample. The vertical line marks Magic Eden's royalty-enforcement change (October 2022). Both series exhibit seasonal variation and co-movement with broader crypto market conditions.

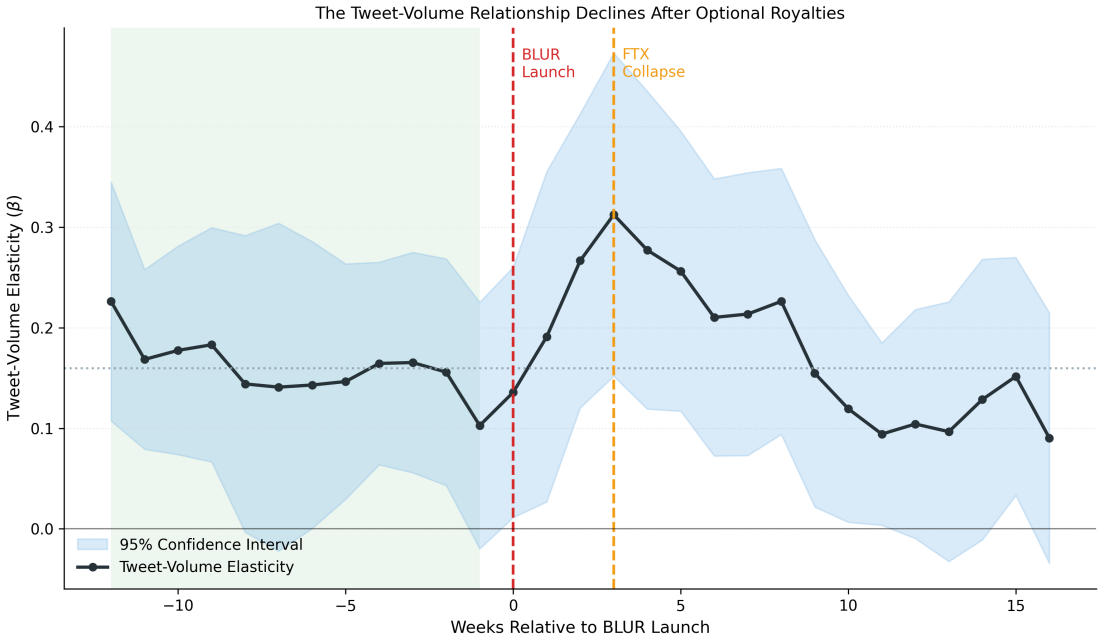

Figure 21: Event Study: Tweet-Volume Elasticity Around Blur Launch (Ethereum)

Notes: Rolling 8-week window estimates of the tweet-volume elasticity from the log–log specification (Appendix C.2), not the baseline semi-elasticity, centered on each week relative to Blur's launch (October 19, 2022, week 0). Shaded region indicates pre-period (weeks −12 to −1). Vertical dashed lines mark Blur launch (red), FTX collapse (orange, November 8, 2022), and OpenSea's capitulation (green, February 17, 2023). 95% confidence intervals shown. The coefficient is approximately stable throughout the pre-period and declines following Blur's launch.

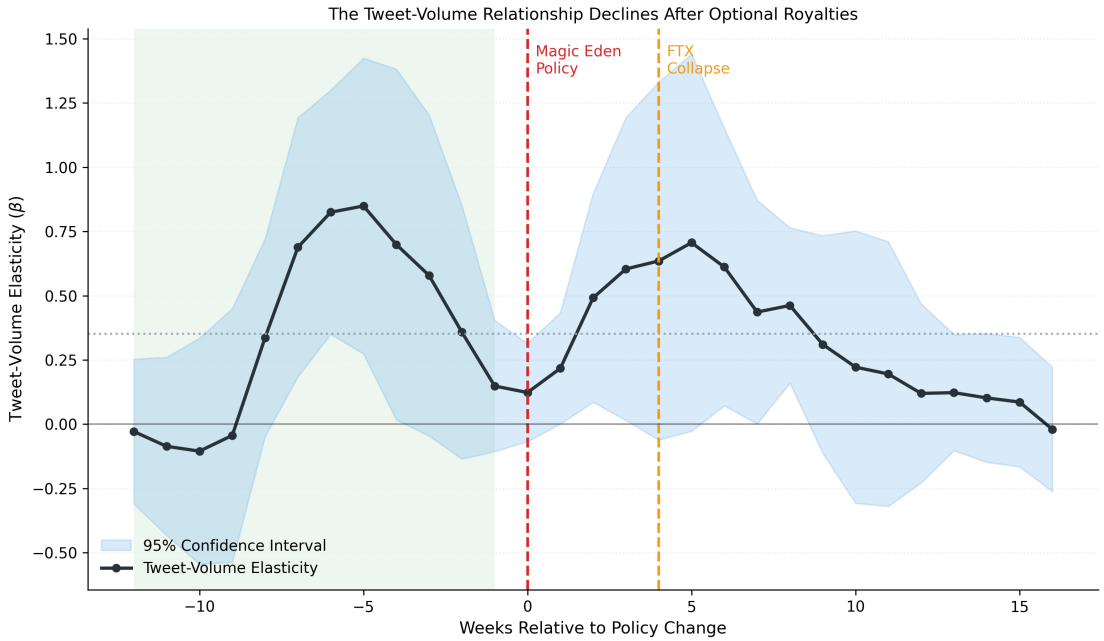

Figure 22: Event Study: Tweet-Volume Elasticity Around Policy Change (Solana)

Notes: Rolling 8-week window estimates of the tweet-volume elasticity from the log–log specification for Solana collections. Week 0 corresponds to Magic Eden's royalty-enforcement change in October 2022. 95% confidence intervals shown. The wider confidence bands relative to Appendix Figure 21 reflect the smaller Solana sample.

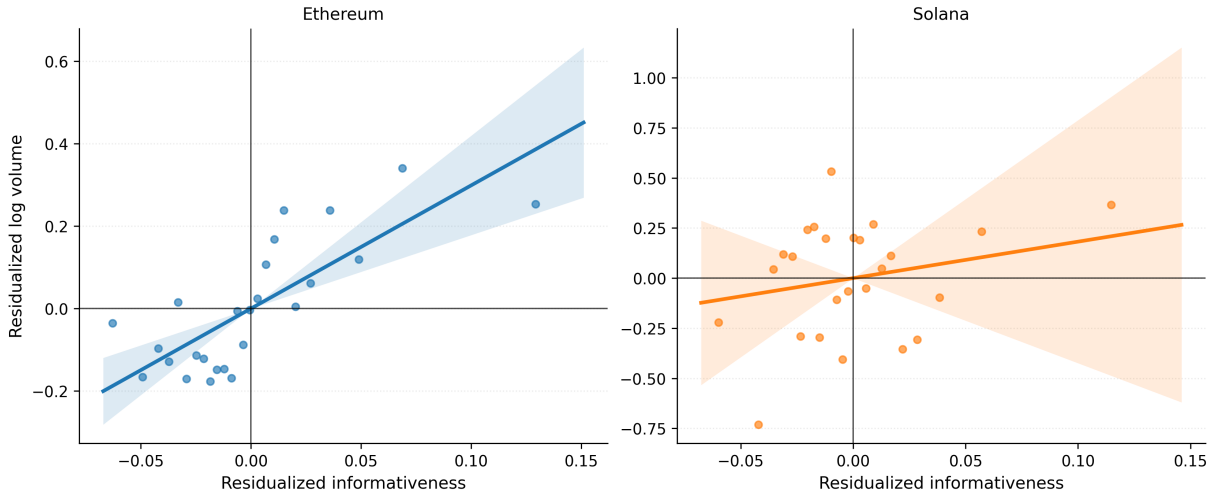

Figure 23: Tweet Informativeness and Trading Volume: Partial Correlations

Notes: Dots are binned means of residualized log volume against residualized tweet informativeness, after partialling out log tweet quantity and collection and week fixed effects (Frisch–Waugh–Lovell). Solid lines plot linear fits with 95% confidence bands based on standard errors two-way clustered by collection and week. The sample restricts to tweet-active collection-weeks (tweet_count > 0). Left: Ethereum. Right: Solana.

Table 19: Information vs. Attention Around the Policy Change (Solana)

	(1)
Log(Announcements + 1)	0.645
	(0.455)
Log(Other tweets + 1)	0.575**
	(0.290)
Log(Announcements + 1) × Post-policy	-0.429
	(0.585)
Log(Other tweets + 1) × Post-policy	-0.243
	(0.388)
(Announcements × Post) − (Other × Post)	-0.186
	(0.883)
Implied Announcements (post)	0.215
	(0.297)
Implied Other (post)	0.332
	(0.228)
Observations	12,450
Collection FE	Yes
Week FE	Yes

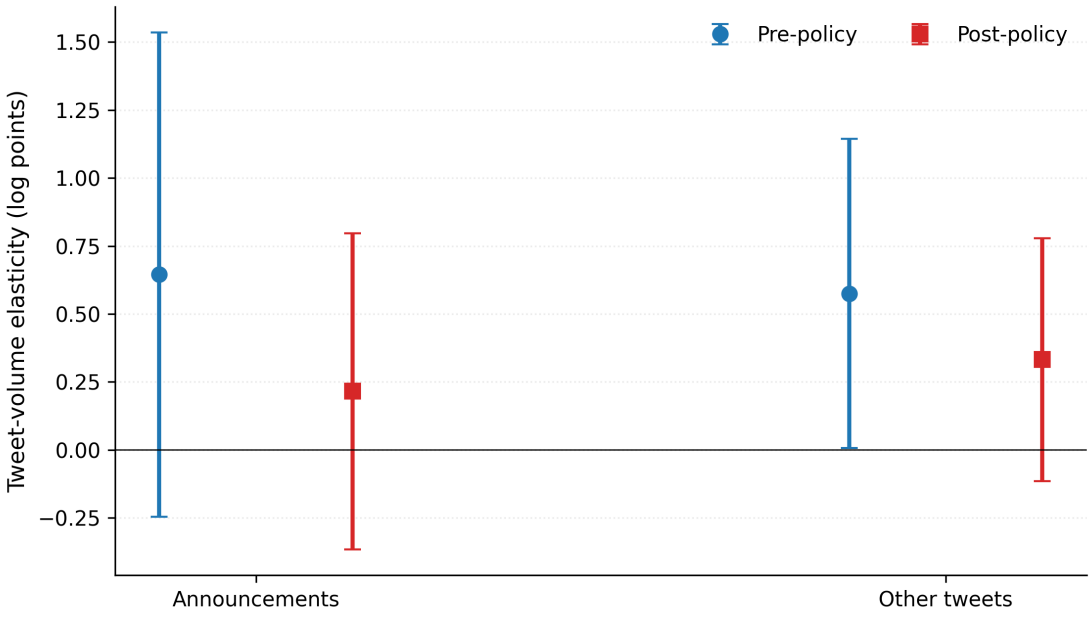

Figure 24: Policy Change and Announcements vs. Other Tweets (Solana)

Notes: Coefficients and 95% confidence intervals implied by Table 19. Standard errors are two-way clustered by collection and week. Week 0 corresponds to Magic Eden's October 2022 royalty-enforcement change (see Section 3.5).

Table 20: Tweet Informativeness and Trading Around the Policy Change (Solana)

	(1)
Log(Tweets + 1)	0.411**
	(0.193)
Informativeness (pre)	7.533
	(8.042)
Informativeness × Post-policy	-7.218
	(9.294)
Informativeness (post)	0.315
	(3.775)
Observations	3,833
Collection FE	Yes
Week FE	Yes

Figure 25: Policy Change and the Informativeness-Volume Link (Solana)

Notes: Coefficients and 95% confidence intervals implied by Appendix Table 20. The sample restricts to collection-weeks with positive tweet activity (tweet_count > 0). All specifications include collection and week fixed effects, with standard errors two-way clustered by collection and week.

Table 21: Channel Decomposition: Transactions vs. Prices (Solana)

	(1) Log(Transactions)	(2) Log(Avg Price)
Log(Tweets + 1)	0.117***	0.130**
	(0.045)	(0.062)
Observations	7,979	7,979
Collection FE	Yes	Yes
Week FE	Yes	Yes

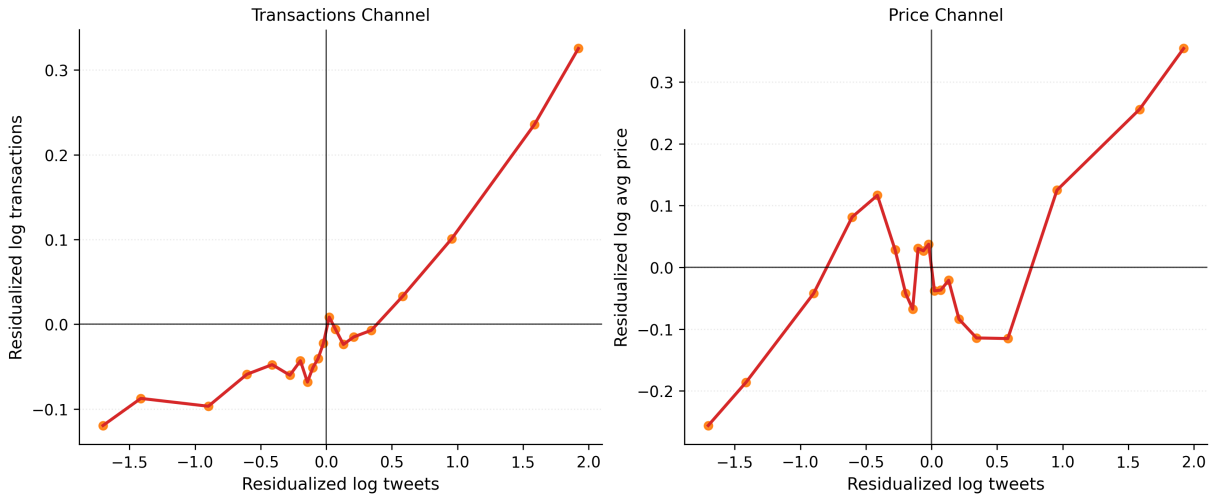

Figure 26: Channel Decomposition: Tweets Are Associated with Both Transactions and Prices (Solana)

Notes: Binned scatters of residualized $\log(1 + \text{Transactions})$ and residualized $\log(1 + \text{Avg Price})$ against residualized $\log(1 + \text{Tweets})$, after partialling out collection and week fixed effects (Frisch–Waugh–Lovell). Lines connect bin means. Standard errors in Table 21 are two-way clustered by collection and week.

Table 22: Fee Heterogeneity Around the Policy Change (Solana)

	(1)
Log(Tweets + 1)	0.719***
	(0.232)
Log(Tweets + 1) × Post-policy	-0.315
	(0.293)
Log(Tweets + 1) × (Fee Rate − Mean, pp)	-1.532*
	(0.843)
Log(Tweets + 1) × Post-policy × (Fee Rate − Mean, pp)	1.580
	(1.120)
Observations	12,450
Collection FE	Yes
Week FE	Yes

Figure 27: Policy Change and Fee Heterogeneity (Solana)

Notes: Implied tweet-volume elasticities at low vs. high contractual fee rates, before and after enforcement weakening. Coefficients are based on a TWFE specification interacting $\log(1 + \text{Tweets})$ with the erosion-period indicator and with fee rates (in percentage points). Standard errors are two-way clustered by collection and week.

Table 23: Placebo Breaks for the Policy-Induced Slope Change (Ethereum and Solana)

	Ethereum	Solana
True $\Delta\hat{\beta}$	-0.013	-0.311
SE(true)	0.086	0.331
p-value (true)	0.878	0.348
Randomization p	0.931	0.491
Placebo mean	-0.020	-0.107
Placebo sd	0.225	0.376
Placebo 5th pct.	-0.390	-0.648
Placebo 95th pct.	0.313	0.416
# placebos	131	106

Table 24: Tweet-Volume Elasticity by ETH Price Regime (Ethereum)

	Low	Mid	High
Log(Tweets + 1)	0.192***	0.173***	0.332***
	(0.039)	(0.051)	(0.106)
Observations	6,563	6,235	4,605
Collection FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes

Notes: The sample is split into terciles based on the weekly average ETH/USD exchange rate. "Low," "Mid," and "High" refer to the bottom, middle, and top terciles. Each subsample is estimated separately with collection and week fixed effects. Standard errors are clustered by collection (the price-regime split absorbs most week-level variation within each subsample). $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

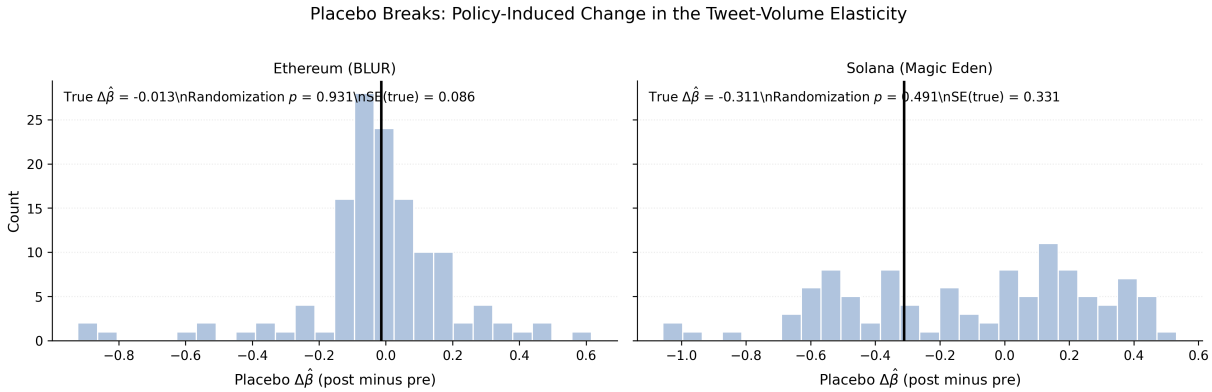

Figure 28: Placebo Break Distribution for the Policy-Induced Slope Change (Ethereum and Solana)

Notes: Histograms show the distribution of placebo pre/post changes in the tweet-volume semi-elasticity when the break date is set to alternative weeks (placebos) rather than the true policy date. For each break date, we estimate the pre- and post-policy slopes in separate TWFE regressions (collection and week fixed effects) within an event-time window of $[-12, +16]$ weeks, and compute the post-minus-pre difference using a stacked cluster-robust covariance that allows within-collection correlation across periods. Placebo candidates exclude the ± 8 weeks around the true policy date and require sufficient observations on both sides of the break. The vertical line marks the estimated change at the true policy break.

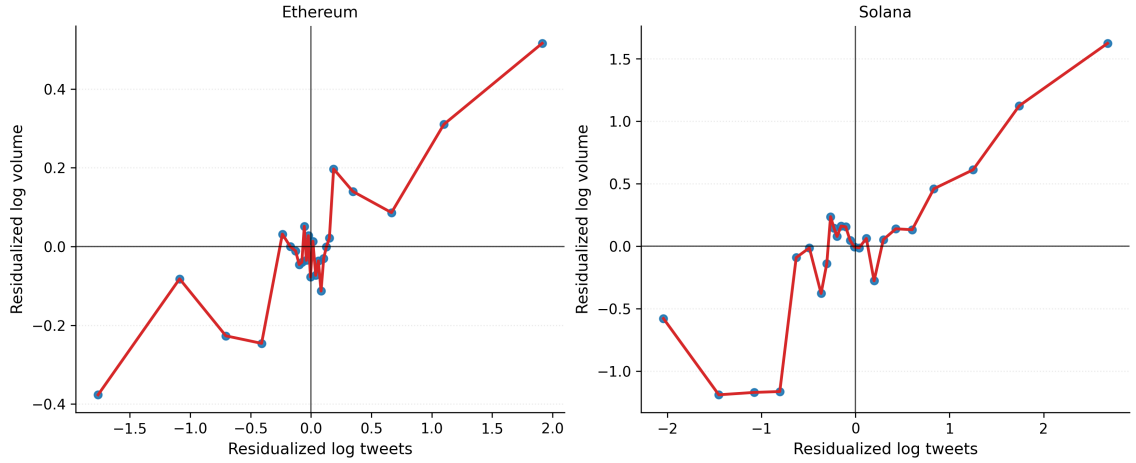

Figure 29: Binned Residual Scatter: Tweets and Trading Volume

Notes: Binned scatter plot of residualized $\log(1 + \text{Volume})$ against residualized $\log(1 + \text{Tweets})$, after partialling out collection and week fixed effects (Frisch–Waugh–Lovell). Dots represent binned means. The solid line is a linear fit. The approximately linear relationship supports the log-linear functional form used throughout.

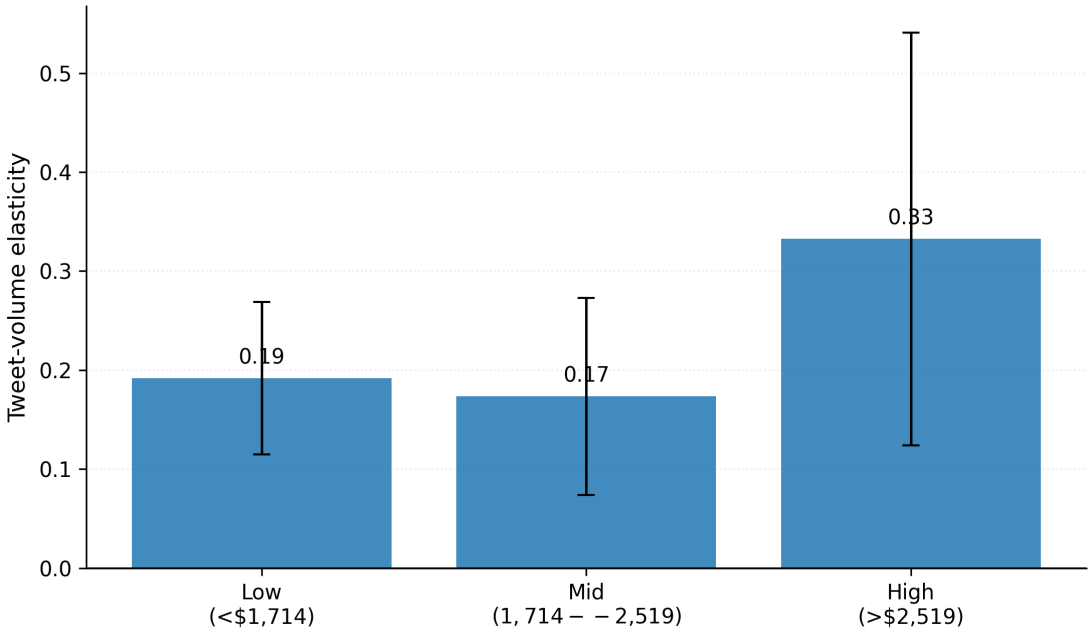

Figure 30: Tweet-Volume Elasticity by ETH Price Regime (Ethereum)

Notes: Coefficients and 95% confidence intervals from Table 24. The tweet-volume semi-elasticity is positive and significant across all ETH price terciles, with a somewhat stronger association during high-price periods. Standard errors are clustered by collection.

Table 25: Robustness Checks (Ethereum)

Specification	Coefficient	Std. Error
B.1: FTX Period Exclusion		
Full Sample	0.225***	(0.061)
Excluding FTX Period	0.227***	(0.065)
B.2: Extreme Observations		
Baseline	0.225***	(0.061)
Winsorized 1%	0.225***	(0.060)
Winsorized 5%	0.207***	(0.052)
B.3: First Differences		
First Differences	0.133***	(0.018)
B.4: Alternative Clustering		
Collection Clustering	0.225***	(0.061)
Two-way Clustering (Collection, Week)	0.225***	(0.061)
Robust (HC1)	0.225***	(0.015)

Table 26: Robustness Checks (Solana)

Specification	Coefficient	Std. Error
B.1: FTX Period Exclusion		
Full Sample	0.580***	(0.165)
Excluding FTX Period	0.602***	(0.174)
B.2: Extreme Observations		
Baseline	0.580***	(0.165)
Winsorized 1%	0.579***	(0.165)
Winsorized 5%	0.571***	(0.163)
B.3: First Differences		
First Differences	0.128**	(0.061)
B.4: Alternative Clustering		
Collection Clustering	0.580***	(0.165)
Two-way Clustering (Collection, Week)	0.580***	(0.165)
Robust (HC1)	0.580***	(0.037)

C.1 The $\log(1+x)$ Transformation: Poisson PML and Split-Sample Analysis

About 70% of Ethereum collection-weeks have zero tweets, so the $\log(1+x)$ transformation produces a mass point at zero. Santos Silva and Tenreyro (2006) show that OLS on log-linearized models can be inconsistent under heteroskedasticity with many zeros. Two checks address this concern.

Poisson pseudo-maximum-likelihood. Poisson PML models the conditional mean of *level* volume directly:

$$\mathbb{E}[\text{Volume}_{it} \mid X] = \exp\big(\alpha_i + \delta_t + \beta \log(1+\text{Tweets}_{it})\big),$$

and is consistent under correct mean specification regardless of the variance function; standard errors are clustered by collection.

Split-sample OLS. We re-estimate the baseline on progressively restricted subsamples: (i) full sample; (ii) tweet-active weeks only; (iii) positive-volume weeks only; (iv) the intersection. If the result were mechanical—$\log(1+0) = 0$ paired on both sides—restricting to the interior would eliminate the coefficient.

Tables 27 and 28 present the results. On Ethereum the OLS coefficient is stable across subsamples and, if anything, rises when restricted to tweet-active weeks (0.277 vs. 0.225 full sample)—zero-tweet weeks contribute noise, not signal. The Poisson PML estimate (0.050, SE = 0.087) is small and statistically insignificant. Because the regressor is $\log(1+\text{Tweets})$, this coefficient is an

elasticity of expected level volume with respect to (1+Tweets), not a semi-elasticity; the Poisson specification weights the right tail of the volume distribution differently from OLS on log volume and is thus not directly numerically comparable. Zero-tweet weeks, not zero-volume weeks, dominate the sample. The Ethereum sample has very few zero-volume weeks, so the Poisson specification offers little additional information beyond the OLS results.

On Solana, where about 36% of collection-weeks have zero volume, the split-sample test is more informative. The Poisson PML coefficient is 0.263 ($p < 0.001$), and restricting to weeks with both positive tweets and positive volume still yields $\hat{\beta} = 0.269$ ($p = 0.012$, $N = 3{,}126$). The association is not an artifact of the log transformation.

Table 27: Poisson PML and Split-Sample Robustness (Ethereum)

Specification	Coefficient	Std. Error	N
Panel A: Poisson PML			
Poisson PML	0.050	(0.087)	17,403
Panel B: OLS Split-Sample			
Full sample (baseline)	0.225***	(0.061)	17,403
Tweeting weeks only (tweet $>$ 0)	0.277***	(0.064)	5,158
Positive-volume weeks only	0.225***	(0.061)	17,403
Both tweet $>$ 0 and volume $>$ 0	0.277***	(0.064)	5,158

Notes: Panel A reports Poisson PML estimates of $\mathbb{E}[\text{Volume}_{it}] = \exp(\alpha_i + \delta_t + \beta \log(1+\text{Tweets}_{it}))$, with standard errors clustered by collection. Panel B re-estimates the baseline OLS TWFE specification (equation 1) on progressively restricted subsamples. All OLS specifications include collection and week fixed effects with two-way clustered standard errors. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Table 28: Poisson PML and Split-Sample Robustness (Solana)

Specification	Coefficient	Std. Error	N
Panel A: Poisson PML			
Poisson PML	0.263***	(0.068)	12,450
Panel B: OLS Split-Sample			
Full sample (baseline)	0.580***	(0.165)	12,450
Tweeting weeks only (tweet $>$ 0)	0.405**	(0.191)	3,833
Positive-volume weeks only	0.248***	(0.089)	8,014
Both tweet $>$ 0 and volume $>$ 0	0.269**	(0.107)	3,126

Notes: Panel A reports Poisson PML estimates with standard errors clustered by collection. Panel B re-estimates the baseline OLS TWFE specification on progressively restricted subsamples. All OLS specifications include collection and week fixed effects with two-way clustered standard errors. The Solana sample has a higher fraction of zero-volume collection-weeks (approximately 36%), making the split-sample comparisons particularly relevant. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

C.2 Log–Log Elasticity Specification

The main text uses a semi-elasticity specification where tweet count enters linearly: $\log(1+\text{Volume}_{it}) = \alpha_i + \delta_t + \beta\,\text{Tweets}_{it} + \varepsilon_{it}$. This yields an intuitive interpretation: one additional tweet is associated with a β% increase in volume. An alternative log–log specification regresses log(1+Volume) on log(1+Tweets), yielding elasticities: a coefficient β means that a 1% increase in tweets is associated with a β% increase in volume. This specification is standard in economics because it is scale-free and works well when both variables vary over many orders of magnitude.

Tables 29 through 32 report log–log estimates for both blockchains. The baseline elasticity on Ethereum (Table 29, column 3) is $\hat{\beta} = 0.225$ ($p < 0.001$), meaning that a 10% increase in tweet activity is associated with a 2.25% increase in trading volume. On Solana (Table 30), the elasticity is more than twice as large: $\hat{\beta} = 0.580$ ($p < 0.001$), consistent with the higher semi-elasticities in the main text.

The policy-change results (Tables 31 and 32) show the same qualitative pattern as the semi-elasticity specifications. On Ethereum, the pre-erosion elasticity is 0.229 and declines to 0.154 after enforcement weakened, a 33% reduction. On Solana, the elasticity drops from 0.59 to 0.52, a 12% reduction. Both specifications—semi-elasticity and log–log elasticity—deliver the same substantive conclusion: the communication–volume link weakened substantially when platforms stopped enforcing creator fees.

The semi-elasticity and log–log specifications are complementary but not numerically comparable. The semi-elasticity measures the percentage change in volume per additional tweet, while the log–log elasticity measures the percentage change in volume per percentage increase in tweets. The marginal effect of one additional tweet differs from the proportional effect of a percentage increase in tweet activity.

Table 29: Log–Log Baseline Specification (Ethereum)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log(Tweets + 1)	0.137*	0.131*	0.225***	0.075***	0.075***	0.133***	0.133***
	(0.077)	(0.074)	(0.061)	(0.015)	(0.015)	(0.019)	(0.019)
Observations	17,403	17,403	17,403	17,234	17,234	17,234	17,234
Collection FE	No	No	Yes	Yes	Yes	No	No
Week FE	No	No	Yes	Yes	Yes	Yes	Yes
Market Controls	No	Yes	No	No	Yes	No	Yes
Specification	Pooled	Pooled	Levels	Levels	Levels	FD	FD

Notes: Dependent variable is log(1+Volume). The regressor is log(1+Tweets), so coefficients are elasticities. All specifications match the main-text Table 3 except for the functional form of the tweet variable. Standard errors are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 30: Log–Log Baseline Specification (Solana)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log(Tweets + 1)	0.881***	0.703***	0.580***	0.114***	0.114***	0.128**	0.128**
	(0.165)	(0.167)	(0.165)	(0.030)	(0.030)	(0.062)	(0.062)
Observations	12,450	12,450	12,450	12,367	12,367	12,367	12,367
Collection FE	No	No	Yes	Yes	Yes	No	No
Week FE	No	No	Yes	Yes	Yes	Yes	Yes
Market Controls	No	Yes	No	No	Yes	No	Yes
Specification	Pooled	Pooled	Levels	Levels	Levels	FD	FD

Notes: Same specification as Table 29 applied to the Solana sample. The elasticity is more than twice as large as on Ethereum. Standard errors are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 31: Log–Log Policy Change (Ethereum)

	Pre-policy	Post-policy
Log(Tweets + 1)	0.229**	0.154***
	(0.105)	(0.040)
Observations	5,738	11,665
Collection FE	Yes	Yes
Week FE	Yes	Yes

Notes: Dependent variable is log(1+Volume). The regressor is log(1+Tweets) interacted with a post-policy dummy. The main coefficient is the pre-erosion elasticity; the interaction captures the change post-erosion. Standard errors are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 32: Log–Log Policy Change (Solana)

	Pre-policy	Post-policy
Log(Tweets + 1)	0.590***	0.520***
	(0.202)	(0.196)
Observations	5,146	7,304
Collection FE	Yes	Yes
Week FE	Yes	Yes

Notes: Same specification as Table 31 applied to the Solana sample. Standard errors are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.3 Local Projection Coefficients

Tables 33 and 34 report all horizon-specific coefficients from the local projections estimated in Section 3.6. The forward direction (tweets → volume) and reverse direction (volume → tweets) are presented side by side for comparison of magnitudes and persistence.

Table 33: Local Projection Coefficients (Ethereum)

	Tweets → Volume			Volume → Tweets		
Horizon h	Coefficient	SE	N	Coefficient	SE	N
0	0.141***	(0.018)	17,065	0.046***	(0.008)	17,065
1	0.101***	(0.022)	16,896	0.033***	(0.008)	16,896
2	0.092***	(0.018)	16,727	0.030***	(0.007)	16,727
3	0.066***	(0.020)	16,558	0.021***	(0.007)	16,558
4	0.086***	(0.019)	16,389	0.016**	(0.007)	16,389
5	0.064***	(0.018)	16,220	0.017**	(0.007)	16,220
6	0.047**	(0.020)	16,052	0.014*	(0.008)	16,052
7	0.041*	(0.022)	15,884	0.025***	(0.008)	15,884
8	0.060**	(0.024)	15,716	0.012*	(0.007)	15,716
Lags	2					
Collection FE	Yes					
Week FE	Yes					
Clustering	Two-way (collection, week)					

Notes: Estimates of $\hat{\beta}_h$ from equation (4) for horizons $h = 0, \ldots, 8$ weeks. "Tweets → Volume" sets $y = \log(1+\text{Volume})$ and $x = \log(1+\text{Tweets})$; "Volume → Tweets" reverses the roles. All regressions include $P = 2$ lags of both variables, collection fixed effects, and week fixed effects. Standard errors in parentheses are two-way clustered by collection and week. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Detailed Disagreement-Model Derivations

This appendix records the algebra behind the posterior formulas, equilibrium price coefficients, disagreement term, disclosure derivatives, and welfare expressions used in Section 4.

D.1 Posterior, Prices, and Disagreement

Effective public-signal precision. Investor i observes

$$\hat{y}_{i,t} = f + \eta_t + \xi_{i,t},$$

Table 34: Local Projection Coefficients (Solana)

	Tweets → Volume			Volume → Tweets		
Horizon h	Coefficient	SE	N	Coefficient	SE	N
0	0.173***	(0.061)	12,284	0.018**	(0.008)	12,284
1	0.216***	(0.059)	12,201	0.019***	(0.007)	12,201
2	0.287***	(0.062)	12,118	0.016**	(0.007)	12,118
3	0.268***	(0.068)	12,035	0.016***	(0.006)	12,035
4	0.264***	(0.068)	11,952	0.008	(0.007)	11,952
5	0.297***	(0.071)	11,869	0.008	(0.007)	11,869
6	0.263***	(0.075)	11,786	0.010	(0.007)	11,786
7	0.254***	(0.074)	11,703	0.002	(0.006)	11,703
8	0.329***	(0.077)	11,620	0.006	(0.006)	11,620
Lags	2					
Collection FE	Yes					
Week FE	Yes					
Clustering	Two-way (collection, week)					

Notes: Estimates of $\hat{\beta}_h$ from equation (4) for horizons $h = 0, \ldots, 8$ weeks. All regressions include $P = 2$ lags of both variables, collection fixed effects, and week fixed effects. Standard errors in parentheses are two-way clustered by collection and week. The Solana sample comprises 83 collections ($N \approx 12{,}300$ at $h = 0$). $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 35: Semi-Elasticity with Collection-Level Effective Fee Interaction

	Ethereum	Solana
Tweet count	0.0073***	0.0157***
	(0.0022)	(0.0047)
Effective fee (demeaned)	24.5170***	−85.6195
	(7.7736)	(64.2225)
Tweet count × Eff. fee	0.0943	1.5811
	(0.1093)	(1.3535)
N	17,403	12,450
R^2	0.0394	0.0147

Notes: Semi-elasticity specification with a continuous effective-fee proxy. The effective fee equals the contractual creator-fee rate (Full_CF_rate) scaled by an enforcement factor: 1 in the pre-erosion regime, 0.5 during partial erosion, and 0.1 under full erosion. The fee variable is demeaned; the interaction Tweets × Effective Fee (dm) tests whether the tweet–volume semi-elasticity varies with the collection-specific fee wedge. All regressions include collection and week fixed effects. Standard errors are two-way clustered by collection and week. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 36: Two-Part (Hurdle) Decomposition: Extensive vs. Intensive Margin (Ethereum)

	(1) Extensive LPM: $\mathbf{1}\{\text{trades} > 0\}$	(2) Intensive OLS: $\log(\text{Vol})\|\text{trades} > 0$	(3) Full Sample OLS: $\log(1 + \text{Vol})$
Tweet count	0.0001	0.0128***	0.0103**
	(0.0000)	(0.0049)	(0.0044)
Tweet count × Post	-0.0000	-0.0085	-0.0047
	(0.0000)	(0.0057)	(0.0050)
N	17,403	9,928	17,403
Collection & Week FE	Yes	Yes	Yes

Notes: Column (1) reports a linear probability model for $\mathbf{1}\{\text{trades} > 0\}$ using the full sample including zero-trade collection-weeks. Column (2) conditions on positive trading activity and estimates the log-volume semi-elasticity (intensive margin). Column (3) reports the full-sample semi-elasticity for comparison. Zero-trade weeks constitute 43.4% of Ethereum collection-week observations (7,624 of 17,552). All regressions include collection and week fixed effects. Standard errors are two-way clustered by collection and week. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 37: Robustness: Excluding Airdrop Windows and Wash-Trading Proxies (Ethereum)

	Tweet count	Tweet × Post	N
Full sample	0.0103**	-0.0047	17,403
	(0.0044)	(0.0050)	
Excl. airdrop window	0.0103**	-0.0044	16,715
	(0.0044)	(0.0051)	
Pre-erosion only	0.0068**	—	5,738
	(0.0031)		
Volume winsorized (99%)	0.0105**	-0.0049	17,403
	(0.0044)	(0.0050)	
Collection & Week FE	Yes	Yes	

Notes: Baseline semi-elasticity specification estimated on four samples. "Full sample" reproduces the baseline. "Excl. airdrop window" drops observations from February 14 to March 14, 2023 (the peak Blur Season 1 airdrop period). "Pre-erosion only" restricts to weeks before October 25, 2022. "Volume winsorized (99%)" clips log volume at the 99th percentile to reduce the influence of extreme wash-trading spikes. All regressions include collection and week fixed effects. Standard errors are two-way clustered by collection and week. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

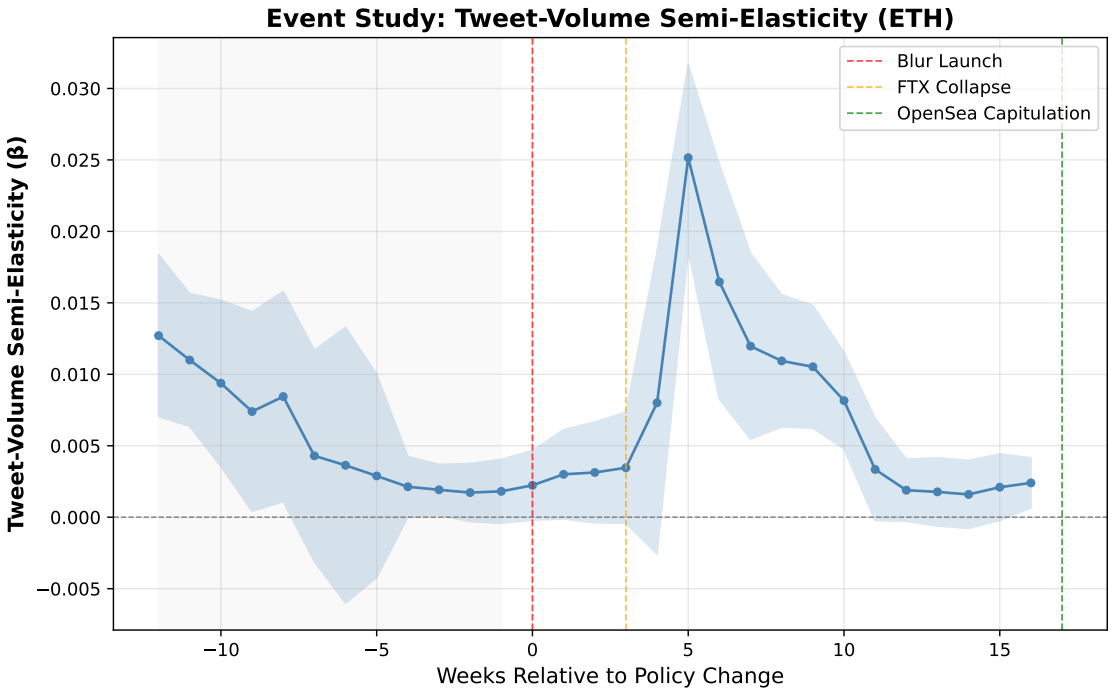

Figure 31: Event Study: Tweet–Volume Semi-Elasticity (Ethereum)

Notes: Rolling 8-week window estimates of the tweet–volume semi-elasticity β from the specification $\log(\text{Volume}_{i,t}) = \alpha_i + \gamma_t + \beta \cdot \text{TweetCount}_{i,t} + \varepsilon_{i,t}$. Each point represents the coefficient from a centered 8-week window at the indicated relative week. Shaded bands show 95% confidence intervals based on heteroskedasticity-robust standard errors. Vertical lines mark the Blur launch (week 0), FTX collapse (week 3), and OpenSea enforcement capitulation (week 17). Unlike Figures 21–22, which plot elasticities from a log-log specification, these estimates use the raw tweet count to match the main semi-elasticity tables.

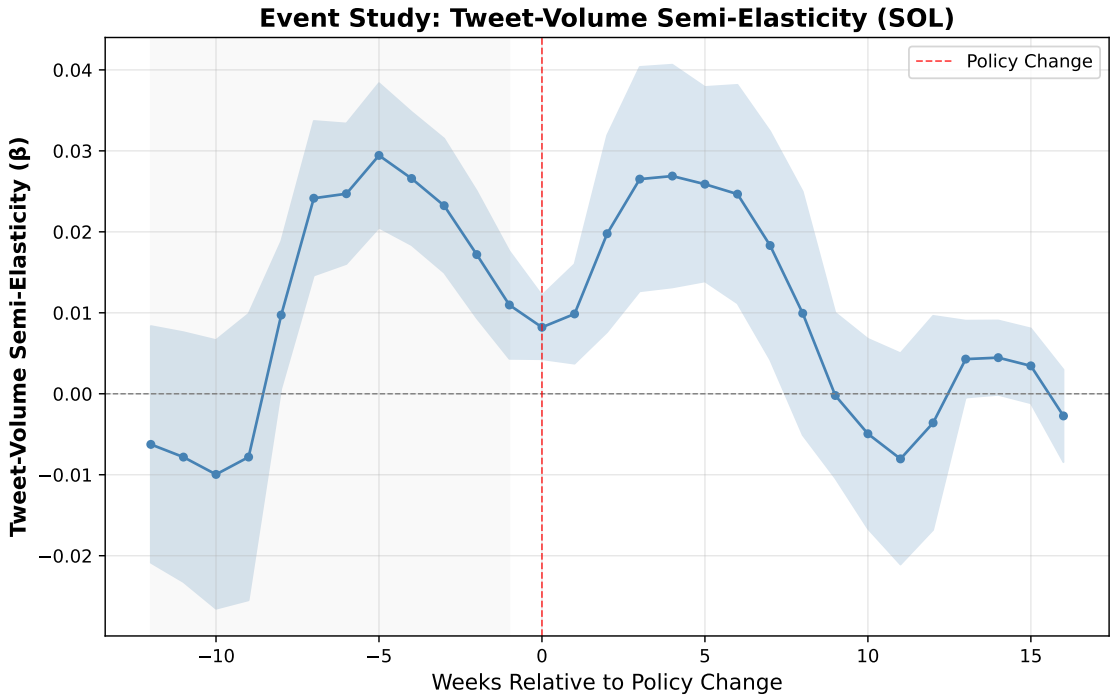

Figure 32: Event Study: Tweet–Volume Semi-Elasticity (Solana)

Notes: Solana counterpart to Figure 31. Rolling 8-week window semi-elasticity estimates around the Solana royalty-enforcement policy change. The semi-elasticity specification uses raw tweet count as the regressor. Shaded bands show 95% confidence intervals. See notes to Figure 31 for details.

where $\eta_t \sim \mathcal{N}(0, \kappa_t^{-1})$ and $\xi_{i,t} \sim \mathcal{N}(0, \psi_t^{-1})$ are independent. Conditional on f, the composite noise term $\eta_t + \xi_{i,t}$ is Normal with variance

$$\mathrm{Var}(\eta_t + \xi_{i,t}) = \kappa_t^{-1} + \psi_t^{-1}.$$

Hence the signal $\hat{y}_{i,t}$ has precision

$$\left(\kappa_t^{-1} + \psi_t^{-1}\right)^{-1} = \frac{\kappa_t \psi_t}{\kappa_t + \psi_t} = \kappa_t^{\text{eff}},$$

which is the harmonic mean reported in the text.

Because the prior and all signals are Gaussian and independent across dates, posterior precision is additive:

$$\sigma_{i,t}^{-2} = \sigma_f^{-2} + \sum_{s=1}^{t} \tau_s + \sum_{s=1}^{t} \kappa_s^{\text{eff}}.$$

The posterior mean is the posterior variance times the precision-weighted sum of the prior mean and all observed signals:

$$\mu_{i,t} = \sigma_{i,t}^2 \left[\sigma_f^{-2}\theta + \sum_{s=1}^{t} \tau_s s_{i,s} + \sum_{s=1}^{t} \kappa_s^{\text{eff}} \hat{y}_{i,s} \right].$$

Substituting $s_{i,s} = f + \varepsilon_{i,s}$ and $\hat{y}_{i,s} = f + \eta_s + \xi_{i,s}$ gives

$$\mu_{i,t} = \sigma_{i,t}^2 \left[\sigma_f^{-2}\theta + \sum_{s=1}^{t} (\tau_s + \kappa_s^{\text{eff}}) f + \sum_{s=1}^{t} \tau_s \varepsilon_{i,s} + \sum_{s=1}^{t} \kappa_s^{\text{eff}} \eta_s + \sum_{s=1}^{t} \kappa_s^{\text{eff}} \xi_{i,s} \right]. \tag{49}$$

Equation (49) separates common components from investor-specific components and is the starting point for the price and volume derivations.

Price coefficients. With CARA–Normal preferences and no per-trade wedge in the pricing benchmark, investor demand is

$$x_{i,t}^{(0)} = \frac{\mu_{i,t} - P_t}{H \sigma_{i,t}^2}.$$

Integrating (49) over the unit mass of investors eliminates the idiosyncratic terms because $\int_0^1 \varepsilon_{i,s}\, di = 0$ and $\int_0^1 \xi_{i,s}\, di = 0$ in the continuum. The average posterior mean is therefore

$$\bar{\mu}_t = \sigma_{i,t}^2 \left[\sigma_f^{-2}\theta + \sum_{s=1}^{t} (\tau_s + \kappa_s^{\text{eff}}) f + \sum_{s=1}^{t} \kappa_s^{\text{eff}} \eta_s \right]. \tag{50}$$

Market clearing,

$$\int_0^1 x_{i,t}^{(0)}\, di + u_t = 1,$$

implies

$$\frac{\bar{\mu}_t - P_t}{H\sigma_{i,t}^2} + u_t = 1 \quad \Longrightarrow \quad P_t = \bar{\mu}_t - H\sigma_{i,t}^2(1-u_t).$$

Substituting (50) yields

$$P_t = \sigma_{i,t}^2\sigma_f^{-2}\theta + \sigma_{i,t}^2\sum_{s=1}^{t}(\tau_s + \kappa_s^{\text{eff}})f + \sigma_{i,t}^2\sum_{s=1}^{t}\kappa_s^{\text{eff}}\eta_s - H\sigma_{i,t}^2 + H\sigma_{i,t}^2 u_t.$$

Matching terms with (13) gives the coefficients in closed form:

$$\begin{aligned}
\alpha_t &= \sigma_{i,t}^2\sigma_f^{-2}\theta - H\sigma_{i,t}^2, \\
\beta_t &= \sigma_{i,t}^2\sum_{s=1}^{t}(\tau_s + \kappa_s^{\text{eff}}), \\
\lambda_{t,s} &= \sigma_{i,t}^2\kappa_s^{\text{eff}}, \qquad s = 1,\ldots,t, \\
\gamma_t &= H\sigma_{i,t}^2.
\end{aligned}$$

Using $\sigma_{i,t}^2\sigma_{i,t}^{-2} = 1$ gives the alternative formula

$$\beta_t = 1 - \sigma_f^{-2}\sigma_{i,t}^2.$$

Taking expectations over $(f, \eta_{\le t}, u_t)$ yields

$$\mathbb{E}[P_t] = \alpha_t + \beta_t\theta = \theta - H\sigma_{i,t}^2.$$

Hence the expected price depends on disclosure only through posterior precision, even though the realized price also loads on public-signal innovations η_s.

Cross-sectional mispricing variance. Subtracting the common price from (49) shows that investor-level mispricing can be written as

$$Z_{i,t} = C_t + \sigma_{i,t}^2\sum_{s=1}^{t}\left(\tau_s\varepsilon_{i,s} + \kappa_s^{\text{eff}}\xi_{i,s}\right), \tag{51}$$

where C_t is common across investors and depends only on $(f, \eta_{\le t}, u_t)$. Conditional on those common states, cross-sectional variation in $Z_{i,t}$ comes only from the idiosyncratic terms. Independence

across dates and signal types therefore gives

$$\begin{aligned}\mathrm{Var}_i(Z_{i,t} \mid f, \eta_{\le t}, u_t) &= \sigma_{i,t}^4 \sum_{s=1}^{t} \left[\tau_s^2 \mathrm{Var}(\varepsilon_{i,s}) + (\kappa_s^{\mathrm{eff}})^2 \mathrm{Var}(\xi_{i,s})\right] \\ &= \sigma_{i,t}^4 \sum_{s=1}^{t} \left[\tau_s + \frac{(\kappa_s^{\mathrm{eff}})^2}{\psi_s}\right],\end{aligned}$$

which is equation (24).

To study the sign of the current-disclosure effect, define

$$A_t := \sum_{s=1}^{t} \tau_s + \sum_{s=1}^{t-1} \frac{(\kappa_s^{\mathrm{eff}})^2}{\psi_s}, \qquad B_t := \sigma_f^{-2} + \sum_{s=1}^{t} \tau_s + \sum_{s=1}^{t-1} \kappa_s^{\mathrm{eff}},$$

and write $k_t := \kappa_t^{\mathrm{eff}}$. Then

$$\sigma_{Z,t}^2 = \frac{A_t + k_t^2/\psi_t}{(B_t + k_t)^2}.$$

Differentiating with respect to k_t gives

$$\frac{\partial \sigma_{Z,t}^2}{\partial k_t} = \frac{2}{(B_t + k_t)^3} \left(\frac{B_t}{\psi_t} k_t - A_t\right). \tag{52}$$

Hence $\sigma_{Z,t}$ is increasing in current effective precision if and only if

$$k_t > \frac{\psi_t A_t}{B_t}. \tag{53}$$

Since

$$\frac{d\kappa_t^{\mathrm{eff}}}{d\kappa_t} = \frac{\psi_t^2}{(\kappa_t + \psi_t)^2} > 0,$$

the same sign condition carries over to raw disclosure κ_t. Equation (53) is the critical-threshold statement used in the main text: sufficiently informative current disclosure widens belief dispersion under disagreement, whereas in the standard common-signal REE the disagreement term disappears and more public information compresses dispersion.

D.2 Trading, Volume, and Disclosure

Trading rules. Conditional on information set $I_{i,t}$, the investor chooses a position x to maximize

$$\mathbb{E}[x(f - P_t) \mid I_{i,t}] - \frac{H}{2}\mathrm{Var}(xf \mid I_{i,t}) = x(\mu_{i,t} - P_t) - \frac{H}{2}\sigma_{i,t}^2 x^2.$$

The first-order condition is

$$\mu_{i,t} - P_t - H\sigma_{i,t}^2 x = 0,$$

so

$$x = \frac{\mu_{i,t} - P_t}{H\sigma_{i,t}^2},$$

which is equation (16).

Date 2. At date 2 the investor solves

$$\max_x \; x(\mu_{i,2} - P_2) - \frac{H}{2}\sigma_{i,2}^2 x^2 - rP_2|x - x_{i,1}|.$$

Let $Z_{i,2} := \mu_{i,2} - P_2$ and $\Delta_2 := rP_2$. On the buy region $x > x_{i,1}$,

$$|x - x_{i,1}| = x - x_{i,1} \quad \Longrightarrow \quad g'(x) = Z_{i,2} - H\sigma_{i,2}^2 x - \Delta_2.$$

Setting $g'(x) = 0$ gives the buy-side target

$$x_{i,2}^+ = \frac{Z_{i,2} - \Delta_2}{H\sigma_{i,2}^2}.$$

On the sell region $x < x_{i,1}$,

$$|x - x_{i,1}| = x_{i,1} - x \quad \Longrightarrow \quad g'(x) = Z_{i,2} - H\sigma_{i,2}^2 x + \Delta_2,$$

which yields the sell-side target

$$x_{i,2}^- = \frac{Z_{i,2} + \Delta_2}{H\sigma_{i,2}^2}.$$

At the kink $x = x_{i,1}$ the subgradient condition is

$$0 \in Z_{i,2} - H\sigma_{i,2}^2 x_{i,1} + \Delta_2[-1, 1],$$

equivalently

$$|Z_{i,2} - H\sigma_{i,2}^2 x_{i,1}| \le \Delta_2.$$

These three cases give Proposition 4.8.

Date 1. At date 1 the optimization problem is

$$\max_x \; G_{i,1}(x) - rP_1|x|, \qquad G_{i,1}(x) := \mathbb{E}[V_{i,2}(x) \mid I_{i,1}].$$

Because $G_{i,1}$ is concave, the no-trade condition at zero is the subgradient condition

$$0 \in \partial G_{i,1}(0) + [-rP_1, rP_1],$$

which yields Proposition 4.9. If $G_{i,1}$ is differentiable at zero, this becomes $|G'_{i,1}(0)| \le rP_1$.

Volume formulas and comparative statics. *Volume representation.* From the date-2 policy,

$$|x^\star_{i,2} - x_{i,1}| = \frac{1}{H\sigma^2_{i,2}}(|\tilde{Z}_{i,2}| - \Delta_2)_+, \qquad \tilde{Z}_{i,2} := \mu_{i,2} - P_2 - H\sigma^2_{i,2}x_{i,1}.$$

Substituting into the definition of dollar volume gives

$$V_2 = P_2 \int_0^1 |x^\star_{i,2} - x_{i,1}|\, di = \frac{P_2}{H\sigma^2_{i,2}} \int_0^1 (|\tilde{Z}_{i,2}| - \Delta_2)_+\, di.$$

Treating the cross-section of investors as a probability space converts the integral into an expectation:

$$V_2 = \frac{P_2}{H\sigma^2_{i,2}}\mathbb{E}[(|\tilde{Z}_{i,2}| - \Delta_2)_+],$$

which is equation (29).

Truncated absolute-Normal expectation. Let $\tilde{Z} \sim \mathcal{N}(\tilde{m}, s^2)$ and define $\Delta \geq 0$. Start from the identity

$$(|\tilde{Z}| - \Delta)_+ = (\tilde{Z} - \Delta)\mathbf{1}\{\tilde{Z} > \Delta\} + (-\tilde{Z} - \Delta)\mathbf{1}\{\tilde{Z} < -\Delta\}.$$

Taking expectations yields

$$\begin{aligned}\mathbb{E}[(|\tilde{Z}| - \Delta)_+] = \mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} > \Delta\}] - \Delta \Pr(\tilde{Z} > \Delta) \\ - \mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} < -\Delta\}] - \Delta \Pr(\tilde{Z} < -\Delta).\end{aligned}$$

Now define

$$a := \frac{\Delta - \tilde{m}}{s}, \qquad b := \frac{\Delta + \tilde{m}}{s}.$$

Standard truncated-Normal identities give

$$\mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} > \Delta\}] = \tilde{m}[1 - \Phi(a)] + s\phi(a),$$

and

$$\mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} < -\Delta\}] = \tilde{m}\Phi(-b) - s\phi(b).$$

Using $\Pr(\tilde{Z} > \Delta) = 1 - \Phi(a)$ and $\Pr(\tilde{Z} < -\Delta) = \Phi(-b) = 1 - \Phi(b)$, substitution gives

$$\mathbb{E}[(|\tilde{Z}| - \Delta)_+] = s[\phi(a) + \phi(b)] + \tilde{m}[\Phi(b) - \Phi(a)] - \Delta[2 - \Phi(a) - \Phi(b)],$$

which is equation (30).

Fee derivative. Holding the frictionless price coefficients fixed, P_2 and $\sigma^2_{i,2}$ are treated as constant

with respect to the fee in the partial-equilibrium exercise. Since $\Delta_2 = rP_2$,

$$\frac{\partial \Delta_2}{\partial r} = P_2.$$

Also,

$$\frac{\partial}{\partial \Delta}(|\tilde{Z}| - \Delta)_+ = -\mathbf{1}\{|\tilde{Z}| > \Delta\}$$

for almost every realization of $\tilde{Z}$. Dominated convergence then gives

$$\frac{\partial}{\partial \Delta_2}\mathbb{E}[(|\tilde{Z}_{i,2}| - \Delta_2)_+] = -\Pr(|\tilde{Z}_{i,2}| > \Delta_2).$$

Applying the chain rule to equation (29),

$$\frac{\partial V_2}{\partial r} = \frac{P_2}{H\sigma_{i,2}^2}\frac{\partial}{\partial \Delta_2}\mathbb{E}[(|\tilde{Z}_{i,2}| - \Delta_2)_+]\frac{\partial \Delta_2}{\partial r} = -\frac{P_2^2}{H\sigma_{i,2}^2}\Pr(|\tilde{Z}_{i,2}| > \Delta_2),$$

which is equation (31).

Disclosure derivative. Because

$$\sigma_{i,t}^{-2} = \sigma_f^{-2} + \sum_{s=1}^{t}\tau_s + \sum_{s=1}^{t}\kappa_s^{\text{eff}},$$

current disclosure enters posterior precision only through κ_t^{eff}. Since

$$\frac{\partial \kappa_t^{\text{eff}}}{\partial \kappa_t} = \frac{\psi_t^2}{(\kappa_t + \psi_t)^2},$$

we obtain

$$\frac{\partial \sigma_{i,t}^{-2}}{\partial \kappa_t} = \frac{\psi_t^2}{(\kappa_t + \psi_t)^2}.$$

Using $d(1/x)/dz = -x^{-2}dx/dz$ with $x = \sigma_{i,t}^{-2}$ gives

$$\frac{\partial \sigma_{i,t}^2}{\partial \kappa_t} = -(\sigma_{i,t}^2)^2\frac{\psi_t^2}{(\kappa_t + \psi_t)^2},$$

which is equation (35). Next write

$$V_t = \frac{P_t}{H\sigma_{i,t}^2}\mathbb{E}[(|\tilde{Z}_{i,t}| - \Delta_t)_+].$$

Under the partial-equilibrium approximation, P_t is held fixed when differentiating with respect to

κ_t, so product differentiation gives

$$\frac{\partial V_t}{\partial \kappa_t} = \frac{P_t}{H\sigma_{i,t}^2}\frac{\partial}{\partial \kappa_t}\mathbb{E}[(|\tilde{Z}_{i,t}| - \Delta_t)_+] + \mathbb{E}[(|\tilde{Z}_{i,t}| - \Delta_t)_+]\frac{\partial}{\partial \kappa_t}\left(\frac{P_t}{H\sigma_{i,t}^2}\right).$$

Since

$$\frac{\partial}{\partial \kappa_t}\left(\frac{1}{\sigma_{i,t}^2}\right) = -\frac{\partial \ln(\sigma_{i,t}^2)}{\partial \kappa_t}\frac{1}{\sigma_{i,t}^2},$$

this becomes

$$\frac{\partial V_t}{\partial \kappa_t} = \frac{P_t}{H\sigma_{i,t}^2}\left[\frac{\partial}{\partial \kappa_t}\mathbb{E}[(|\tilde{Z}_{i,t}| - \Delta_t)_+] - \frac{\partial \ln(H\sigma_{i,t}^2)}{\partial \kappa_t}\mathbb{E}[(|\tilde{Z}_{i,t}| - \Delta_t)_+]\right],$$

which is equation (36). The sign of the first term is governed by the threshold condition in (53).

Royalty revenue and the disclosure first-order condition. *Founder objective.* Under royalties, the founder chooses (κ_1, κ_2) to maximize total expected royalty revenue

$$\Pi_F(\kappa_1, \kappa_2) = r(V_1 + V_2), \qquad \kappa_1 + \kappa_2 = \bar{\kappa}.$$

Defining period-t revenue as $R_t := rV_t$ gives

$$\Pi_F = R_1 + R_2,$$

which is equation (37) aggregated across periods.

Interior first-order condition. Use the resource constraint to write $\kappa_2 = \bar{\kappa} - \kappa_1$. Then

$$\frac{d\Pi_F}{d\kappa_1} = \frac{\partial R_1}{\partial \kappa_1} + \frac{\partial R_2}{\partial \kappa_2}\frac{d\kappa_2}{d\kappa_1} = \frac{\partial R_1}{\partial \kappa_1} - \frac{\partial R_2}{\partial \kappa_2}.$$

At an interior optimum this derivative must equal zero, so

$$\frac{\partial R_1}{\partial \kappa_1} = \frac{\partial R_2}{\partial \kappa_2},$$

which is equation (38). If the left-hand side exceeds the right-hand side, the founder shifts disclosure toward period 1; if the right-hand side is larger, the founder shifts disclosure toward period 2.

Elasticity objects. Because $R_t = rV_t$, the local marginal revenue from disclosure can be written as

$$\frac{\partial R_t}{\partial \kappa_t} = r\frac{\partial V_t}{\partial \kappa_t} = \frac{R_t}{\kappa_t}\omega_t, \qquad \omega_t := \frac{\partial \ln V_t}{\partial \ln \kappa_t}.$$

Hence the ratio of marginal revenues is

$$\frac{\partial R_2/\partial \kappa_2}{\partial R_1/\partial \kappa_1} = \frac{R_2}{R_1} \cdot \frac{\omega_2}{\omega_1} \cdot \frac{\kappa_1}{\kappa_2}.$$

The objects Ψ and Θ in the text isolate the first two ingredients: Ψ measures the relative size of date-2 trading, while $\Theta = \omega_1/\omega_2$ measures the relative responsiveness of volume to disclosure. Larger date-2 revenue or larger date-2 elasticity tilts the founder toward back-loading.

Membership disclosure derivatives. *Counterfactual and revenue equivalence.* Define the benchmark price index

$$\bar{P} := w_1 P_1 + w_2 P_2, \qquad w_1, w_2 \geq 0,\ w_1 + w_2 = 1.$$

The counterfactual membership fee is the deterministic quantity $M := m\,\mathbb{E}[\bar{P}]$. Revenue equivalence chooses m so that the deterministic membership fee equals expected royalty revenue:

$$M = m\,\mathbb{E}[\bar{P}] = \mathbb{E}[r(V_1 + V_2)].$$

The benchmark policy experiment corresponds to the upfront contract $(w_1, w_2) = (1, 0)$, so $M = m\,\mathbb{E}[P_1]$.

Expected prices and disclosure. Expected price at date t is $\mathbb{E}[P_t] = \theta - H\sigma_{i,t}^2$, so disclosure incentives under membership follow from how disclosure changes posterior variance. Using

$$\frac{d\sigma_{i,t}^2}{d\kappa_j} = -(\sigma_{i,t}^2)^2 \frac{d\sigma_{i,t}^{-2}}{d\kappa_j},$$

the date-1 derivative is

$$\frac{\partial \mathbb{E}[P_1]}{\partial \kappa_1} = H(\sigma_{i,1}^2)^2 \frac{\psi_1^2}{(\kappa_1 + \psi_1)^2} > 0.$$

Along the budget constraint $\kappa_2 = \bar{\kappa} - \kappa_1$, the date-2 derivative is

$$\frac{d\mathbb{E}[P_2]}{d\kappa_1} = H(\sigma_{i,2}^2)^2 \left[\frac{\psi_1^2}{(\kappa_1 + \psi_1)^2} - \frac{\psi_2^2}{(\kappa_2 + \psi_2)^2}\right].$$

For a general weighted membership contract $\bar{P} = w_1 P_1 + w_2 P_2$, the founder's derivative is therefore

$$\frac{d\mathbb{E}[\bar{P}]}{d\kappa_1} = w_1 \frac{\partial \mathbb{E}[P_1]}{\partial \kappa_1} + w_2 \frac{d\mathbb{E}[P_2]}{d\kappa_1}.$$

Front-loading under the benchmark contract. Under the benchmark upfront contract $(w_1, w_2) = (1, 0)$, the derivative is strictly positive everywhere, so the solution is the front-loading corner. Once $w_2 > 0$, the date-2 term can offset the date-1 gain, so interior disclosure splits are possible. The paper's sharp front-loading theorem is stated for the upfront membership benchmark.

D.3 Welfare

Royalties. For a given trading rule $(x^\star_{i,1}, x^\star_{i,2})$, investor wealth net of transfers equals expected payoff from the NFT position minus the CARA risk penalty minus the proportional fees paid when positions are changed. Taking expectations gives

$$W(r) = \mathbb{E}[f] - \frac{H}{2}\mathbb{E}\big[\sigma^2_{i,2}(x^\star_{i,2})^2\big] - rP_2\mathbb{E}[|x^\star_{i,2} - x^\star_{i,1}|] - rP_1\mathbb{E}[|x^\star_{i,1}|],$$

which is equation (41). The first term is the expected payoff, the second is the certainty-equivalent risk cost, and the final two terms are expected royalty payments.

State-by-state wedge loss. Fix one period, define Z as the frictionless mispricing, and let $\Delta = rP$. The investor solves

$$U_\Delta(Z) := \max_x \left\{ Zx - \frac{H\sigma^2}{2}x^2 - \Delta|x| \right\}.$$

Without the wedge, the first-order condition $Z - H\sigma^2 x = 0$ gives

$$x^0 = \frac{Z}{H\sigma^2}, \qquad U_0(Z) = \frac{Z^2}{2H\sigma^2}.$$

With the wedge, the optimizer is

$$x^\star = \begin{cases} 0, & |Z| \le \Delta, \\ \dfrac{Z - \operatorname{sgn}(Z)\Delta}{H\sigma^2}, & |Z| > \Delta. \end{cases}$$

Evaluating the objective at $x^\star$ gives

$$U_\Delta(Z) = \begin{cases} 0, & |Z| \le \Delta, \\ \dfrac{(|Z| - \Delta)^2}{2H\sigma^2}, & |Z| > \Delta. \end{cases}$$

Subtracting from $U_0(Z)$ yields

$$U_0(Z) - U_\Delta(Z) = \frac{1}{2H\sigma^2}\left[Z^2\mathbf{1}\{|Z| \le \Delta\} + \left(2|Z|\Delta - \Delta^2\right)\mathbf{1}\{|Z| > \Delta\}\right].$$

Taking expectations gives equation (42). The expression is strictly positive for any $\Delta > 0$ and any continuously distributed Z with nonzero support.

Revenue-equivalent membership. Under membership, the per-trade wedge is zero, so investors attain the frictionless surplus $U_0(Z)$ in each period and pay the fixed fee M once. Under royalties, they attain $U_\Delta(Z)$ and additionally pay the proportional transfer $\Delta|x^\star|$. At revenue equivalence, expected transfers are the same across regimes, so the welfare difference reduces to the difference in net trading surplus:

$$W(M) - W(r) = \sum_{t=1}^{2} \mathbb{E}[U_0(Z_t) - U_{\Delta_t}(Z_t) - \Delta_t|x_t^\star|].$$

Using the piecewise form of $x^\star$,

$$U_0(Z) - U_\Delta(Z) - \Delta|x^\star| = \frac{1}{2H\sigma^2}\left[Z^2\mathbf{1}\{|Z| \le \Delta\} + \Delta^2\mathbf{1}\{|Z| > \Delta\}\right],$$

which is strictly positive almost surely when Z has a continuous distribution. Summing across dates gives the welfare ranking in Proposition 4.14. Membership removes the wedge that suppresses information-based rebalancing, not just a transfer from the trading margin.

E Supplementary Proofs

Proof of Proposition 4.4. Start from the averaged posterior mean in equation (50):

$$\bar{\mu}_t = \sigma_{i,t}^2\left[\sigma_f^{-2}\theta + \sum_{s=1}^{t}(\tau_s + \kappa_s^{\text{eff}})f + \sum_{s=1}^{t}\kappa_s^{\text{eff}}\eta_s\right].$$

Market clearing implies

$$P_t = \bar{\mu}_t - H\sigma_{i,t}^2(1 - u_t).$$

Substituting for $\bar{\mu}_t$ gives

$$P_t = \sigma_{i,t}^2\sigma_f^{-2}\theta + \sigma_{i,t}^2\sum_{s=1}^{t}(\tau_s + \kappa_s^{\text{eff}})f + \sigma_{i,t}^2\sum_{s=1}^{t}\kappa_s^{\text{eff}}\eta_s - H\sigma_{i,t}^2 + H\sigma_{i,t}^2 u_t.$$

The constant term is therefore

$$\alpha_t = \sigma_{i,t}^2\sigma_f^{-2}\theta - H\sigma_{i,t}^2.$$

Since

$$\beta_t = \sigma_{i,t}^2 \sum_{s=1}^{t} (\tau_s + \kappa_s^{\text{eff}}) = 1 - \sigma_f^{-2}\sigma_{i,t}^2,$$

we have $\sigma_{i,t}^2 \sigma_f^{-2} = 1 - \beta_t$, so

$$\alpha_t = (1 - \beta_t)\theta - H\sigma_{i,t}^2.$$

Taking unconditional expectations and using $\mathbb{E}[f] = \theta$, $\mathbb{E}[\eta_s] = 0$, and $\mathbb{E}[u_t] = 0$ gives

$$\mathbb{E}[P_t] = \alpha_t + \beta_t\theta = \theta - H\sigma_{i,t}^2.$$

□

Proof of Theorem 4.6. Under Assumption 4.3, posterior precision $\sigma_{i,t}^{-2}$ is pinned down directly by primitives:

$$\sigma_{i,t}^{-2} = \sigma_f^{-2} + \sum_{s=1}^{t} (\tau_s + \kappa_s^{\text{eff}}).$$

Given $\sigma_{i,t}^{-2}$, the pricing coefficients are explicit:

$$\beta_t = \frac{\sum_{s=1}^{t} (\tau_s + \kappa_s^{\text{eff}})}{\sigma_{i,t}^{-2}}, \qquad \lambda_{t,s} = \frac{\kappa_s^{\text{eff}}}{\sigma_{i,t}^{-2}}, \qquad \gamma_t = H\sigma_{i,t}^2, \qquad \alpha_t = (1 - \beta_t)\theta - H\sigma_{i,t}^2.$$

There is therefore no fixed-point problem left in the coefficients once the primitives are given. Any two candidate equilibria must share the same posterior precision and hence the same explicit coefficient formulas, so they coincide. Existence follows by direct substitution of these coefficients into the price ansatz, and uniqueness follows because the formulas determine a single coefficient vector. □

Proof of Proposition 4.8. Let $Z_{i,2} := \mu_{i,2} - P_2$ and $\Delta_2 := rP_2$. The date-2 objective is

$$g(x) := xZ_{i,2} - \frac{H\sigma_{i,2}^2}{2}x^2 - \Delta_2|x - x_{i,1}|.$$

The function is strictly concave, so the optimum is characterized by piecewise first-order conditions.

If $x > x_{i,1}$, then $|x - x_{i,1}| = x - x_{i,1}$ and

$$g'(x) = Z_{i,2} - H\sigma_{i,2}^2 x - \Delta_2.$$

Setting $g'(x) = 0$ gives

$$x = \frac{Z_{i,2} - \Delta_2}{H\sigma_{i,2}^2} = x_{i,2}^{+}.$$

This candidate is feasible in the $x > x_{i,1}$ region if and only if

$$Z_{i,2} - H\sigma_{i,2}^2 x_{i,1} > \Delta_2.$$

If $x < x_{i,1}$, then $|x - x_{i,1}| = x_{i,1} - x$ and

$$g'(x) = Z_{i,2} - H\sigma_{i,2}^2 x + \Delta_2.$$

Setting $g'(x) = 0$ gives

$$x = \frac{Z_{i,2} + \Delta_2}{H\sigma_{i,2}^2} = x_{i,2}^-,$$

which is feasible if and only if

$$Z_{i,2} - H\sigma_{i,2}^2 x_{i,1} < -\Delta_2.$$

At the kink $x = x_{i,1}$, the subgradient condition is

$$0 \in Z_{i,2} - H\sigma_{i,2}^2 x_{i,1} + \Delta_2[-1, 1],$$

which is equivalent to

$$|Z_{i,2} - H\sigma_{i,2}^2 x_{i,1}| \leq \Delta_2.$$

These three regions exhaust the state space and yield the policy in Proposition 4.8. □

Proof of Proposition 4.9. Define the expected continuation value at date 1 as $G_{i,1}(x) := \mathbb{E}[V_{i,2}(x) \mid I_{i,1}]$. In the date-2 problem (25), the objective is concave in $(x_{i,2}, x_{i,1})$ because it is a quadratic concave function of $x_{i,2}$ plus the term $-rP_2|x_{i,2} - x_{i,1}|$, which is concave in $(x_{i,2}, x_{i,1})$. Maximizing a concave function over $x_{i,2} \in \mathbb{R}$ preserves concavity in the parameter $x_{i,1}$, so $V_{i,2}(x_{i,1})$ is concave in $x_{i,1}$. Taking conditional expectations preserves concavity, hence $G_{i,1}$ is concave.

The date-1 objective in (26) can be written as

$$\Pi_{i,1}(x) = G_{i,1}(x) - rP_1|x|,$$

which is concave. A necessary and sufficient optimality condition for $x^\star = 0$ in concave maximization is $0 \in \partial\Pi_{i,1}(0)$. Using $\partial|x| \big|_{x=0} = [-1, 1]$, we have

$$\partial\Pi_{i,1}(0) = \partial G_{i,1}(0) - rP_1\, \partial|x| \big|_{x=0} = \partial G_{i,1}(0) + [-rP_1,\ rP_1],$$

which yields the stated condition. If $G_{i,1}$ is differentiable at 0, then $\partial G_{i,1}(0) = \{G'_{i,1}(0)\}$ and the condition reduces to $|G'_{i,1}(0)| \leq rP_1$.

Finally, the no-trade region expands in r because the interval $[-rP_1, rP_1]$ is increasing in r. □

Proof of Proposition 4.10. From Proposition 4.8, the optimal date-2 trade size satisfies

$$|x^\star_{i,2} - x_{i,1}| = \frac{1}{H\sigma^2_{i,2}} \left(|\tilde{Z}_{i,2}| - \Delta_2\right)_+, \qquad \tilde{Z}_{i,2} := \mu_{i,2} - P_2 - H\sigma^2_{i,2} x_{i,1},\ \Delta_2 := rP_2.$$

Substituting into the dollar-volume definition (27) and integrating across the unit mass of investors gives (29). Under heterogeneous interpretation, the cross-sectional distribution of $\tilde{Z}_{i,2}$ has variance given by equation (24), which includes the cumulative disagreement components $\sum_{s=1}^{2} (\kappa_s^{\text{eff}})^2/\psi_s$. This wider dispersion (relative to the standard model) is the mechanism through which disagreement raises trading volume.

For the closed form, let $\tilde{Z} \sim \mathcal{N}(\tilde{m}, s^2)$ and $\Delta \geq 0$. Since $(|\tilde{Z}| - \Delta)_+ = (\tilde{Z} - \Delta)\mathbf{1}\{\tilde{Z} > \Delta\} + (-\tilde{Z} - \Delta)\mathbf{1}\{\tilde{Z} < -\Delta\}$, we have

$$\mathbb{E}[(|\tilde{Z}| - \Delta)_+] = \mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} > \Delta\}] - \Delta \Pr(\tilde{Z} > \Delta) - \mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} < -\Delta\}] - \Delta \Pr(\tilde{Z} < -\Delta).$$

Define $a := (\Delta - \tilde{m})/s$ and $b := (\Delta + \tilde{m})/s$. A change of variables $u = (z - \tilde{m})/s$ gives the truncated moment identities

$$\begin{aligned}
\mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} > \Delta\}] &= \int_\Delta^\infty z \frac{1}{s}\phi\left(\frac{z - \tilde{m}}{s}\right) dz \\
&= \tilde{m} \int_a^\infty \phi(u)\, du + s \int_a^\infty u\phi(u)\, du \\
&= \tilde{m}[1 - \Phi(a)] + s\phi(a), \\
\mathbb{E}[\tilde{Z}\mathbf{1}\{\tilde{Z} < -\Delta\}] &= \int_{-\infty}^{-\Delta} z \frac{1}{s}\phi\left(\frac{z - \tilde{m}}{s}\right) dz \\
&= \tilde{m} \int_{-\infty}^{-b} \phi(u)\, du + s \int_{-\infty}^{-b} u\phi(u)\, du \\
&= \tilde{m}\Phi(-b) - s\phi(b),
\end{aligned}$$

where we used $\int_a^\infty u\phi(u)\, du = \phi(a)$, $\int_{-\infty}^c u\phi(u)\, du = -\phi(c)$, and $\phi(-b) = \phi(b)$. Also $\Pr(\tilde{Z} > \Delta) = 1 - \Phi(a)$ and $\Pr(\tilde{Z} < -\Delta) = \Phi(-b) = 1 - \Phi(b)$. Substituting these into the display above yields (30). □

Proof of Proposition 4.12. Fix Z and consider the wedge-distorted certainty-equivalent objective

$$U_\Delta(Z) := \max_{x \in \mathbb{R}} \left\{ Zx - \frac{H\sigma^2}{2} x^2 - \Delta|x| \right\}.$$

When $\Delta = 0$, the first-order condition $Z - H\sigma^2 x = 0$ gives

$$x^0 = \frac{Z}{H\sigma^2}, \qquad U_0(Z) = \frac{Z^2}{2H\sigma^2}.$$

When $\Delta > 0$, the argument from Proposition 4.8 with $x_{i,1} = 0$ yields

$$x^\star = \begin{cases} 0, & |Z| \le \Delta, \\ \dfrac{Z - \operatorname{sgn}(Z)\Delta}{H\sigma^2}, & |Z| > \Delta. \end{cases}$$

Evaluating the objective at this optimizer gives

$$U_\Delta(Z) = \begin{cases} 0, & |Z| \le \Delta, \\ \dfrac{(|Z| - \Delta)^2}{2H\sigma^2}, & |Z| > \Delta. \end{cases}$$

Subtracting $U_\Delta(Z)$ from $U_0(Z)$ yields

$$U_0(Z) - U_\Delta(Z) = \frac{1}{2H\sigma^2}\Big[Z^2 \mathbf{1}\{|Z| \le \Delta\} + \big(2|Z|\Delta - \Delta^2\big)\mathbf{1}\{|Z| > \Delta\}\Big].$$

Taking expectations gives equation (42). Strict positivity follows because the integrand is nonnegative everywhere and strictly positive on a set of positive probability whenever $\Delta > 0$ and Z has a continuous distribution with nonzero support. □

Proof of Proposition 4.13. Under the benchmark upfront contract $(w_1, w_2) = (1, 0)$, founder revenue is proportional to $\mathbb{E}[P_1]$. Since

$$\mathbb{E}[P_1] = \theta - H\sigma_{i,1}^2,$$

the derivative with respect to disclosure is

$$\frac{\partial \mathbb{E}[P_1]}{\partial \kappa_1} = -H \frac{\partial \sigma_{i,1}^2}{\partial \kappa_1}.$$

From

$$\sigma_{i,1}^{-2} = \sigma_f^{-2} + \tau_1 + \kappa_1^{\text{eff}}, \qquad \kappa_1^{\text{eff}} = \frac{\kappa_1 \psi_1}{\kappa_1 + \psi_1},$$

we have

$$\frac{\partial \sigma_{i,1}^{-2}}{\partial \kappa_1} = \frac{\psi_1^2}{(\kappa_1 + \psi_1)^2},$$

and therefore

$$\frac{\partial \sigma_{i,1}^2}{\partial \kappa_1} = -(\sigma_{i,1}^2)^2 \frac{\psi_1^2}{(\kappa_1 + \psi_1)^2}.$$

Substituting back gives

$$\frac{\partial \mathbb{E}[P_1]}{\partial \kappa_1} = H(\sigma_{i,1}^2)^2 \frac{\psi_1^2}{(\kappa_1 + \psi_1)^2} > 0$$

for all $\kappa_1 \in [0, \bar{\kappa}]$. Hence the objective is strictly increasing in κ_1, so the unique optimum is the

corner solution $\kappa_1^\star = \bar{\kappa}$ and $\kappa_2^\star = 0$. □

Proof of Proposition 4.14. Fix a period and let Z denote the frictionless mispricing and $\Delta = rP$ the proportional wedge. Under royalties, the investor attains the certainty-equivalent surplus

$$U_\Delta(Z) = \max_x \left\{ Zx - \frac{H\sigma^2}{2}x^2 - \Delta|x| \right\},$$

and pays the proportional transfer $\Delta|x^\star|$, where $x^\star$ is the wedge-distorted optimizer. Under membership, the per-trade wedge is zero, so the investor attains $U_0(Z)$ and pays the fixed fee M once.

At revenue equivalence, aggregate expected transfers are identical across the two regimes, so the welfare difference reduces to

$$U_0(Z) - U_\Delta(Z) - \Delta|x^\star|.$$

Using the piecewise optimizer

$$x^\star = \begin{cases} 0, & |Z| \le \Delta, \\ \dfrac{Z - \operatorname{sgn}(Z)\Delta}{H\sigma^2}, & |Z| > \Delta, \end{cases}$$

we obtain

$$U_0(Z) - U_\Delta(Z) - \Delta|x^\star| = \frac{1}{2H\sigma^2}\Big[Z^2\mathbf{1}\{|Z| \le \Delta\} + \Delta^2\mathbf{1}\{|Z| > \Delta\}\Big].$$

This expression is strictly positive almost surely when Z has a continuous distribution and $\Pr(|Z| > 0) > 0$. Taking expectations and summing across the two trading dates gives $W(M) > W(r)$. □