



UNIVERSITÉ DU QUÉBEC EN OUTAOUAIS
DÉPARTEMENT DES SCIENCES ADMINISTRATIVES

WORKING PAPER NO. 15

Faster, Wider, More Fragile? The Evolution of Intraday Shock Propagation in U.S. Equity Markets, 1992–2026

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

July 2026

Version 1.0

Abstract

We study how intraday shocks travel across U.S. equities using 7.69 billion one-minute bars covering the full listed and delisted cross-section from 1992 to 2026. Three facts emerge. First, markets absorb idiosyncratic shocks: across 1.85 million isolated $\pm 6\sigma$ one-minute events, abnormal market-wide propagation—measured against matched within-stock placebos—is economically zero at horizons of 1 to 120 minutes, in every decade. Second, transmission is concentrated and has intensified: the shocked stock’s twenty closest lagged-correlation neighbors respond strongly ($+64.0 \times 10^{-4}$ abnormal jump probability at five minutes), and this network channel tripled between the 1990s and the 2020s while market-wide spillover stayed at zero. Third, tail synchronization migrated to the market-wide minute: ordinary extreme minute returns became roughly 108 times less frequent than in the 1990s, the Epps curve flattened (one-minute pairwise correlation 0.004 in 1996 versus 0.30 in 2015), and severe co-jump minutes spike by orders of magnitude in 2008 and 2020. Modern markets absorb the idiosyncratic and transmit the systemic.

Keywords: intraday shock propagation, market fragility, high-frequency data, contagion, co-jumps, Epps effect, market microstructure

JEL Classification: G10, G12, G14

1. Introduction

When one security suffers a violent price move, does the shock stay local or does it spread? The answer shapes almost everything that happens inside a trading day: how fast prices aggregate information, how much liquidity providers must charge for standing in the market, how execution algorithms should schedule parent orders, and whether an idiosyncratic event can grow into a market-wide dislocation. Regulators confront the same question from the other side: episodes such as May 6, 2010 showed that modern equity markets can move from normal functioning to synchronized collapse within minutes, faster than any human intervention. Yet the empirical record on intraday shock transmission is built almost entirely from short samples, small panels, or single episodes. Whether the propagation mechanism itself has changed over the three decades in which trading moved from exchange floors to fragmented, algorithmic, ETF-intermediated electronic markets is, to our knowledge, undocumented.

The tension motivating this paper is that theory predicts change in both directions. Faster, cheaper, more competitive markets should absorb ordinary idiosyncratic shocks more efficiently: liquidity provision reacts within seconds, arbitrage capital is abundant in normal times, and information is incorporated nearly continuously. But the same institutions that supply this efficiency—high-speed market making with tight inventory limits, cross-asset arbitrage, index and basket intermediation—withdraw or transmit pressure precisely when many securities move together. Modern markets may therefore absorb ordinary shocks more efficiently while transmitting extreme shocks more rapidly and discontinuously.

We ask how the speed, breadth, persistence, and asymmetry of intraday shock propagation across U.S. equities evolved over the past three decades, and under which market conditions an initially idiosyncratic shock becomes sectoral or market-wide.

Our laboratory is a panel of 7,686,275,524 one-minute OHLCV bars, originally sourced from Finnhub.io, covering 81,306 tickers—including delisted securities—from January 1992 through March 2026. The data are audited exhaustively before use: observed trading dates match the official exchange calendar exactly over thirty-four years, timestamp conventions are verified on both sides of every daylight-saving transition, and integrity violations are negligible (100 impossible OHLC bars in 7.69 billion). Because the source provides bars only for minutes with trades, and identifies securities by ticker alone, the empirical design works with liquidity-screened, lagged-rank universes, ticker spells rather than tickers, and detected corporate-action windows; Section 3 documents every screen.

Three findings organize the paper. First, *markets absorb isolated idiosyncratic shocks—and we can measure how completely*. We detect 1.85 million one-minute events in which a single top-500 stock moves by at least six robust standard deviations of its own residual return while the rest of the

market is calm. Relative to matched within-stock placebos, the abnormal share of other securities responding within 5 to 120 minutes is economically zero (point estimates of -1 to -3×10^{-4} with day-block bootstrap standard errors of about 1×10^{-4}), in every decade, for positive and negative shocks alike. Idiosyncratic events die where they land.

Second, *the transmission of common information has become nearly instantaneous*. In 1996 the average pairwise correlation of one-minute returns among the hundred largest stocks was 0.004—economically zero, the classical Epps effect—while thirty-minute returns already comoved at 0.16. By 2015 the one-minute correlation was 0.30 against 0.39 at thirty minutes: the Epps curve has flattened almost completely. Information that once took half an hour to spread across the cross-section now spreads within the minute.

Third, *tail risk has migrated from the security to the minute*. Ordinary extreme moves became rare: the frequency of one-minute returns exceeding 1% fell from roughly 518 per 10,000 in the 1990s to 4.8 per 10,000 in the 2010s—a factor of about 108. But synchronized extremes did not: minutes in which more than 2% of the universe jumps beyond 4σ simultaneously—systemic co-jump minutes—spike by orders of magnitude in 2008 and 2020, and the frequency of $|r| > 5\%$ minute returns rose from 0.015 per 10,000 in 2017–2019 to 5.8 in 2020, a factor of 391. The joint pattern—absorption of the idiosyncratic, synchronization of the systemic—is the modern form of market fragility.

Identification receives explicit attention throughout. The central threat is that securities respond to common information rather than to each other. Our shock definition residualizes against a leave-one-out market factor with strictly lagged loadings and scales; isolation filters require a calm market and fewer than 2% simultaneous extremes at the event minute; every event is compared to a placebo in the same stock at the same minute of day on a calm day; and the pre-event window shows no abnormal response, ruling out anticipation or mechanical selection. What the isolation filters remove—the simultaneous-extreme minutes—we do not discard but study separately as the systemic co-jump sample. Common-information contamination is thus measured, not assumed away.

Mechanically, we ask *where* propagation should hide if it exists: responses concentrated among the shocked stock’s closest lagged-correlation neighbors, cumulative rather than point-in-time responses, and state dependence across calm and stressed markets and across shock magnitudes. Section 7 reports these probes.

The paper contributes to four literatures. To the contagion literature, which measures transmission at daily or lower frequency for small panels, we contribute event-level evidence from the full U.S. cross-section at the minute scale over thirty-four years. To the market-microstructure and high-frequency-trading debate, which contrasts normal-times liquidity gains with episodic fragility, we contribute the missing time series: propagation and synchronization measured uniformly across

the entire structural transformation of U.S. trading. To the co-jump econometrics literature, we contribute a long-sample account of how systemic co-jump frequency and severity evolved. And to the measurement literature, we document what OHLCV bars can and cannot identify about liquidity and transmission, with every limitation quantified.

Section 3 describes and audits the data. Section 4 defines shocks, outcomes, and networks. Section 5 presents the empirical design. Section 6 reports the descriptive transformation of the market and the baseline propagation results. Section 7 examines networks, cumulative responses, and co-jumps. Section 8 summarizes robustness, and Section 9 concludes.

2. Related literature

This paper measures how quickly, how broadly, how persistently, and how asymmetrically idiosyncratic intraday shocks travel across the U.S. equity cross-section, and how that transmission evolved between 1992 and 2026. The question sits at the intersection of several literatures: theories of cross-asset shock transmission; empirical work on fire sales, flows, and ownership linkages; the measurement of comovement, commonality, and liquidity dynamics; lead–lag effects and information diffusion; network approaches to contagion and systemic risk; the transformation of U.S. market structure and the high-frequency trading debate; the econometrics of jumps, co-jumps, and scheduled announcements; and the measurement toolkit for high-frequency and OHLCV data. We take each in turn and close with the paper’s positioning.

2.1. Theories of cross-asset shock transmission

Theory identifies two broad families of channels through which a shock to one security moves others. The first is informational. [King and Wadhwani \(1990\)](#) model rational agents who infer information from *other* markets’ price changes, so that an idiosyncratic “mistake” in one market transmits broadly— a mechanism designed to explain why nearly all markets crashed together in October 1987. [Kodres and Pritsker \(2002\)](#) show that cross-market portfolio rebalancing of shared macro risk exposures transmits idiosyncratic shocks even without common fundamentals, with severity increasing in information asymmetry. The second family operates through balance sheets and funding. [Kyle and Xiong \(2001\)](#) generate contagion as a wealth effect: convergence traders who lose money in one market liquidate positions in others, raising volatility and correlations exactly when capital is scarce. [Yuan \(2005\)](#) shows that borrowing constraints make crises, contagion, and asymmetric price movements likelier after intermediate losses. [Brunnermeier and Pedersen \(2009\)](#) link market liquidity to funding liquidity in mutually reinforcing margin spirals that explain sudden, correlated liquidity dry-ups and flight to quality. Most directly related to our intraday setting, [Cespa and Foucault \(2014\)](#) show that when liquidity providers learn from *other securities’ prices*,

illiquidity becomes contagious and the market admits liquidity-crash equilibria—a mechanism they connect explicitly to May 6, 2010. These theories jointly predict that transmission should be state-dependent (stronger when volatility is high and capital constrained), asymmetric (stronger for negative shocks), and increasing in the degree to which pricing and liquidity provision are cross-security activities. Our measures of propagation speed, breadth, asymmetry, and market-state dependence are designed to test exactly these comparative statics at the minute scale.

2.2. Fire sales, flows, and ownership networks

A complementary empirical literature documents non-fundamental transmission through investor portfolios. Mutual funds facing outflows engage in forced sales concentrated in their holdings, generating price pressure and reversals (Coval and Stafford, 2007); predictable flows generate predictable price pressure (Lou, 2012); and uninformative trading imbalances in one stock spill over to correlated stocks through risk-averse liquidity providers (Andrade et al., 2008). Ownership structure itself predicts fragility: stocks with concentrated or correlated-flow ownership have higher volatility and comovement (Greenwood and Thesmar, 2011), and common active ownership forecasts cross-sectional correlation beyond fundamentals-based similarity (Antón and Polk, 2014). At the institutional level, overlapping portfolios create fire-sale spillovers that aggregate into measurable systemic vulnerability (Greenwood et al., 2015; Duarte and Eisenbach, 2021). This work establishes that idiosyncratic demand shocks *do* propagate, but its resolution is coarse—monthly or quarterly holdings and flows. Whether the same propagation operates within the hour, and whether its speed has changed as intermediation became algorithmic, is the margin our data address.

2.3. Comovement, commonality, and liquidity dynamics

The commonality literature establishes that liquidity has a systematic component: spreads and depths co-move with market- and industry-wide liquidity (Chordia et al., 2000), common factors in order flow explain much of the commonality in returns (Hasbrouck and Seppi, 2001), and commonality is supplied in part by common intermediaries—stocks sharing a NYSE specialist firm co-move in liquidity beyond the market factor (Coughenour and Saad, 2004). Time variation is systematic as well: liquidity falls after market declines, especially when funding is tight (Hameed et al., 2010); specialist inventories and trading losses forecast spread widening (Comerton-Forde et al., 2010); and returns to liquidity provision spike with the VIX as liquidity “evaporates” in turmoil (Nagel, 2012). Aggregate liquidity risk is priced (Pástor and Stambaugh, 2003; Acharya and Pedersen, 2005), and commonality strengthens in downturns across forty markets (Karolyi et al., 2012). Relative to this literature, we do not measure liquidity *levels*; we measure the *transmission of extreme events* and

use liquidity states—constructed from OHLCV proxies—as conditioning variables. At very short horizons, measured comovement is attenuated mechanically by asynchronous trading (Epps, 1979), a bias any minute-level design must engineer around; our breadth measures therefore condition on trading eligibility rather than assuming synchronous observation.

2.4. Lead–lag effects and gradual information diffusion

Cross-autocorrelation in returns is among the oldest facts in empirical asset pricing: large stocks lead small stocks (Lo and MacKinlay, 1990), predominantly within industries, consistent with gradual information diffusion (Hou, 2007). Economic links sharpen the pattern: returns of principal customers predict suppliers’ returns (Cohen and Frazzini, 2008), and cross-predictability follows supplier–customer industry links broadly, declining with analyst coverage and institutional ownership (Menzly and Ozbas, 2010). At high frequency, tick-data studies document that liquid assets lead illiquid ones with intraday-seasonal lead–lag structure (Huth and Abergel, 2014), and Chincio et al. (2019) show that one-minute-ahead cross-stock predictability exists but is sparse, short-lived, and news-related. These studies measure *unconditional* predictive structure. Our object is different: the conditional response of the cross-section to an identified extreme idiosyncratic event, its horizon profile, and the secular evolution of both. The lead–lag literature nonetheless supplies one of our falsification benchmarks: propagation estimated through lagged-return networks should reproduce known liquidity-ordered asymmetries if it is measuring genuine transmission.

2.5. Networks: real linkages, financial linkages, and measured connectedness

Network theory shows when idiosyncratic shocks fail to diversify away: sufficiently asymmetric input–output structures let micro shocks aggregate (Acemoglu et al., 2012), and intermediate integration and diversification maximize cascade susceptibility in financial cross-holdings (Elliott et al., 2014); Glasserman and Young (2016) survey the contagion-network literature. Real-economy links propagate firm-level shocks measurably—natural-disaster shocks to suppliers impose large value losses on customers (Barrot and Sauvagnat, 2016)—and network positions carry asset-pricing implications (Herskovic, 2018). On the measurement side, connectedness is estimated from variance decompositions (Diebold and Yilmaz, 2012, 2014), Granger-causality graphs among financial institutions (Billio et al., 2012), sparse VARs (Barigozzi and Brownlees, 2019), and tail-risk networks (Hautsch et al., 2015). The contagion-measurement tradition warns that unconditional correlation comparisons are biased in volatile regimes (Forbes and Rigobon, 2002); defining contagion as comovement beyond fundamentals, Bekaert et al. (2014) find “wake-up call” transmission tied to weak fundamentals. Classic volatility-spillover evidence—meteor showers in FX (Engle et al., 1990), international open/close spillovers (Hamao et al., 1990), and cross-market volatility

transmission (King and Wadhwani, 1990)—anticipates our design at coarser resolution and for a handful of markets. Across all of this work the networks are estimated at daily or lower frequency on small panels. We estimate pre-shock networks on minute data for thousands of securities and use them as lagged conditioning information—asking whether network position predicts transmitters, receivers, amplifiers, and absorbers—with permuted-network placebos in the spirit of the falsification tradition.

2.6. Market-structure evolution and the high-frequency trading debate

Our sample spans essentially the entire modern transformation of U.S. equity trading, and the literature documents each step. Decimalization in 2001 cut spreads, particularly for active stocks (Bessembinder, 2003); liquidity itself promotes price efficiency (Chordia et al., 2008); turnover exploded in the 2000s alongside algorithmic trading (Chordia et al., 2011); the NYSE’s 2006–07 Hybrid automation made prices more efficient while raising the cost of immediacy (Hendershott and Moulton, 2011); and fragmentation across venues did not measurably harm quality (O’Hara and Ye, 2011). On algorithmic and high-frequency trading, the normal-times evidence is broadly positive: algorithmic trading narrows spreads and improves price discovery (Hendershott et al., 2011; Brogaard et al., 2014), modern HFT market makers halve spreads on entry (Menkveld, 2013), low-latency activity improves short-horizon quality (Hasbrouck and Saar, 2013), machines improve FX efficiency (Chaboud et al., 2014), and colocation-based evidence replicates internationally, with higher short-term volatility as the caveat (Boehmer et al., 2021); Menkveld (2016) surveys the channel-by-channel evidence. Theory tempers the optimism: speed investment can be socially excessive (Biais et al., 2015), fast news trading concentrates volume and short-run price correlation (Foucault et al., 2016), toxic arbitrage degrades liquidity (Foucault et al., 2017), the continuous limit order book hard-wires a latency arms race (Budish et al., 2015)—measured directly in message data as races won in microseconds, a fifth of volume (Aquilina et al., 2022)—and speed can crowd out information acquisition (Weller, 2018; Baldauf and Mollner, 2020), with HFTs eventually trading *with* informed institutional orders (van Kervel and Menkveld, 2019). O’Hara (2015) argues the machine age changes what microstructure concepts mean. This literature motivates our central hypothesis—normal-times absorption improving while tail transmission accelerates—but it measures liquidity levels, price discovery, or single events. We contribute the missing time series: propagation speed and breadth of extreme shocks, measured uniformly across three decades of structural change, with flexible year effects rather than imposed trend breaks.

2.7. Fragility and flash events

Episodic evidence shows modern liquidity can vanish abruptly. On May 6, 2010, high-frequency traders did not trigger the crash but amplified selling once it began (Kirilenko et al., 2017); cross-market arbitrage broke down, leaving the seller dependent on local liquidity (Menkveld and Yueshen, 2019); and fragmentation shaped which ETFs were hit (Madhavan, 2012). More systematically, HFTs supply liquidity in single-stock extreme moves but their provision weakens when *many* stocks move together (Brogaard et al., 2018)—precisely the margin our breadth measure tracks. Proposed early-warning metrics remain contested: VPIN signaled deteriorating conditions before the Flash Crash (Easley et al., 2012), but its power appears mechanically tied to volatility and trading intensity (Andersen and Bondarenko, 2014). This debate cautions against fragility indicators that repackage volatility; our propagation measures are therefore defined relative to lagged conditional distributions and validated against matched placebo times.

2.8. ETFs, indexing, and passive intermediation

A structural change specific to the later sample is the growth of index-linked intermediation. ETF ownership raises non-fundamental volatility of underlying stocks through the arbitrage channel (Ben-David et al., 2018) and contributes to return comovement, especially in small and illiquid names (Da and Shive, 2018); index-product arbitrage flipped index return serial dependence from positive to negative across twenty markets (Baltussen et al., 2019). Effects on efficiency are contested: ETF ownership is associated with weaker firm-specific information environments (Israeli et al., 2017), yet ETF activity speeds the incorporation of systematic earnings information where information environments are weak (Glosten et al., 2021), and Russell-reconstitution evidence suggests price efficiency survives passive growth because the remaining informed investors adjust (Coles et al., 2022). For us, the ETF layer is both a candidate mechanism for the *widening* of propagation (basket arbitrage converts idiosyncratic shocks into correlated flows) and a required robustness screen: no headline result may be driven by ETFs themselves.

2.9. Jumps, co-jumps, and scheduled announcements

Our shock measurement builds directly on high-frequency econometrics. Intraday volatility periodicity must be filtered before anything else (Andersen and Bollerslev, 1997); bipower variation separates jumps from diffusion (Barndorff-Nielsen and Shephard, 2004); formal jump tests exist at the observation level (Lee and Mykland, 2008; Aït-Sahalia and Jacod, 2009). Two cautions discipline our thresholds: at ultra-high frequency much of what coarser methods call jumps is friction or “drift bursts” (Christensen et al., 2014), and multiple-testing corrections eliminate a large share of spurious detections, with true jumps rarer and news-related (Bajgrowicz et al., 2016).

The co-jump literature is the nearest empirical neighbor to our propagation object: index-level jumps are systematic co-jumps linked to macro news while idiosyncratic jumps diversify away (Bollerslev et al., 2008); simultaneous multivariate jumps are frequent, associated with liquidity dry-ups, and priced (Caporin et al., 2017); jumps coincide with liquidity-demand shocks and news (Boudt and Petitjean, 2014); and simultaneous jumps across countries carry portfolio consequences (Das and Uppal, 2004). Self- and cross-excitation of jumps has a canonical language in Hawkes processes (Hawkes, 1971; Bacry et al., 2015), estimated on international indices by Aït-Sahalia et al. (2015). Scheduled macro announcements provide the identified common shocks against which idiosyncratic propagation must be benchmarked: announcement surprises produce jump-like conditional-mean responses in FX and global markets (Andersen et al., 2003, 2007), generate the cojumps documented by Lahaye et al. (2011), and structure expected returns around announcement times (Savor and Wilson, 2013; Lucca and Moench, 2015). Accordingly, our baseline excludes announcement windows and treats announcement-window propagation as a separate, labeled object. Where the co-jump literature tests for *simultaneity*, our measures trace the *sequence*—breadth and intensity at one through 120 minutes—and its secular change.

2.10. High-frequency measurement: factor structure, OHLCV proxies, and machine learning

Separating idiosyncratic from systematic movement at the minute scale requires a factor structure: common factors exist in high-frequency returns and order flow (Hasbrouck and Seppi, 2001), high-frequency PCA is consistent with low-frequency covariance structure (Aït-Sahalia and Xiu, 2019), latent continuous and jump factors are estimable and economically interpretable (Pelger, 2020), and betas decompose into continuous, jump, and overnight components with distinct pricing (Bollerslev et al., 2016); signed semicovariances refine the decomposition (Bollerslev et al., 2020). Jump-tail risk itself commands a large, time-varying premium (Bollerslev and Todorov, 2011), which is why our propagation outcomes separate tail responses from ordinary ones. Because our data are bars rather than quotes, liquidity enters only through proxies with a long pedigree: spread estimators from return autocovariance (Roll, 1984) and high–low ranges (Corwin and Schultz, 2012), range-based volatility (Parkinson, 1980; Garman and Klass, 1980), price-impact ratios (Amihud, 2002), and trade-based information measures (Hasbrouck, 1991; Goyenko et al. (2009) validate such proxies against TAQ benchmarks. For estimation we use local projections (Jordà, 2005) for horizon-specific responses, and the machine-learning extension follows the out-of-sample discipline of Gu et al. (2020), informed by evidence that microstructure features carry universal predictive content (Sirignano and Cont, 2019) and remain informative in the machine age (Easley et al., 2021).

2.11. Positioning and contribution

Each literature above contributes a piece: theory predicts state-dependent, asymmetric transmission; the flows literature demonstrates non-fundamental propagation at monthly resolution; the network literature measures connectedness on small panels at daily frequency; the market-structure literature documents the transformation of trading but measures liquidity levels and single events; the co-jump literature tests simultaneity at high frequency over short samples. To our knowledge—based on a documented search dated 2026-07-13 (replication package)—no prior study combines (i) minute-level resolution, (ii) the full U.S. equity cross-section including delisted securities, and (iii) a 1992–2026 sample spanning floor trading, decimalization, Reg NMS fragmentation, and ETF-intermediated, algorithmic markets. The joint analysis of propagation speed, breadth, persistence, asymmetry, secular evolution, and network conditioning at that resolution and length—with placebo-timing, permuted-network, and announcement-exclusion falsifications—is the contribution.

3. Data

3.1. Source and structure

Our data are one-minute OHLCV bars for U.S.-listed securities originally sourced from Finnhub.io (<https://finnhub.io>), a market data provider, and aggregated from the provider’s monthly archives into 411 monthly Parquet files covering January 1992 through March 2026 (87.7 GB compressed) with no missing months. We access the aggregation through a public archive whose exact byte-level snapshot is pinned and checksum-verified in the replication package, fixing the data used throughout for reproducibility. Each row records the open, high, low, and close price (U.S. dollars, unadjusted for corporate actions), share volume, and ticker for one minute of trading in one security. Timestamps are UTC and denote the start of the bar minute; we verify the convention empirically by mapping bar times to the Eastern-time trading session in months on both sides of daylight-saving transitions.

A bar exists only for minutes with at least one reported trade: there are no zero-volume rows in the source. Missing minutes are therefore informative about trading activity, and we treat them explicitly rather than filling a regular grid.

3.2. Pilot audit

Before downloading the archive we audited three pilot months (June 1992, June 2009, December 2025; 58.8 million bars). Trading-day counts match the U.S. exchange calendar in all three months, including the December 24, 2025 early close. No OHLC-consistency violations ($L_{i,t} > \min(O_{i,t}, C_{i,t})$, $H_{i,t} < \max(O_{i,t}, C_{i,t})$, or $L_{i,t} > H_{i,t}$), null values, negative volumes, or zero-volume




Figure 1. Dataset coverage. Panel (a): compressed size of each monthly source file, 1992–2026 (all 411 files of the pinned snapshot). Panel (b): number of symbols with at least one bar per trading day in the three pilot months.

bars occur in the pilot sample; 14 duplicate ticker–minute pairs (all in inactive over-the-counter names) appear in December 2025. Liquid securities are densely covered throughout: the four most active tickers in December 2025 have complete 390-bar regular sessions, and large issuers already average 277–369 bars per day in June 1992.

Coverage widens over time, from roughly 8,300 symbols per month in 1992 to about 23,100 in late 2025; much of the recent expansion reflects over-the-counter instruments rather than listed common equity, and our sample filters hold the universe definition fixed using lagged dollar-volume ranks. Only 22–26% of June 1992 tickers appear in the later pilot months, indicating that delisted securities are retained in the historical files.

Figure 1 summarizes raw data volume and per-day symbol counts; Figure 2 documents the familiar intraday patterns in absolute returns, volume, and range. Table 1 reports pilot summary statistics for the 500 highest dollar-volume securities in each pilot month.

3.3. Full-sample audit and clean panel

The full archive contains 7,686,275,524 bars. Observed trading dates match the official XNYS calendar exactly—zero missing sessions and zero dates outside the calendar over thirty-four years—and the regular-hours share of bars is flat across calendar months, ruling out any daylight-saving mishandling. Integrity failures are negligible: 100 impossible OHLC bars (concentrated in 1993–1995) and 10,128 duplicate ticker-minute keys (1.3 per million), which we merge under a logged deterministic rule. The panel spans 81,306 distinct tickers forming 221,957 continuous spells; 29,112 tickers have multiple spells, so spells, not tickers, are the identity unit. Entries and exits occur continuously (delisted securities are present throughout), with one discrete source-coverage




Figure 2. Intraday seasonality, pilot months. Minute-of-day means during regular trading hours for the top 500 securities by within-month dollar volume: (a) absolute one-minute log return in basis points; (b) share of the stock-day’s volume; (c) high–low log range in basis points. Descriptive full-month normalization; no corporate-action screens.

Table 1. Summary statistics, pilot months

Month	Variable	Mean	SD	p1	p5	p50	p95	p99	N
June 1992	Minute log return (bps)	-0.17	255.04	-145.99	-66.01	0.00	63.90	145.28	1,374,563
	High–low log range (bps)	12.57	230.25	0.00	0.00	0.00	63.09	155.85	1,385,486
	Dollar volume per bar (\$000)	140.01	824.08	1.86	3.06	31.32	579.15	1,561.88	1,385,486
	Valid RTH bars per stock-day	126.84	79.58	19.00	35.00	107.00	303.00	369.00	10,923
June 2009	Minute log return (bps)	-0.16	1,172.84	-38.30	-19.51	0.00	19.18	37.47	4,190,657
	High–low log range (bps)	28.67	945.13	0.00	0.00	9.49	36.04	65.47	4,201,656
	Dollar volume per bar (\$000)	886.22	19,849.91	9.21	28.84	285.18	2,807.64	8,382.52	4,201,656
	Valid RTH bars per stock-day	382.00	31.56	270.00	357.00	389.00	390.00	390.00	10,999
December 2025	Minute log return (bps)	-0.02	10.72	-32.15	-14.12	0.00	14.02	31.78	3,904,759
	High–low log range (bps)	8.33	12.25	0.00	0.00	4.66	28.94	57.64	3,915,747
	Dollar volume per bar (\$000)	2,833.57	13,640.57	58.50	142.61	846.13	8,849.93	37,017.92	3,915,747
	Valid RTH bars per stock-day	356.37	53.42	166.00	223.00	383.00	390.00	390.00	10,988

Notes: Regular trading hours (09:30–16:00 ET, exchange calendar) for the top 500 securities by within-month dollar volume in each pilot month of the pinned Finnhub.io minute-bar aggregation. Minute log returns use consecutive same-day closes; the high–low range is $\log(H) - \log(L)$ per one-minute bar; bars exist only for minutes with at least one reported trade. Prices are unadjusted for splits; the pilot is descriptive and imposes no corporate-action screens.

expansion in 2020 that our fixed-rule universes absorb (Figure 3). A conservative overnight-ratio screen detects 6,944 corporate-action events (4,819 splits, 2,125 reverse splits); event days ± 1 trading day are excluded from all return-based analyses. Table 2 tracks the sample from raw bars to the regular-hours analysis panel.

3.4. Known limitations

The archive contains no quotes, spreads, depth, or trade direction; all liquidity statements in this paper refer to OHLCV-based proxies. Securities are identified by ticker only, prices are unadjusted, and security types are unlabeled; the corporate-action detection, ticker-spell construction, and universe classification underlying the final sample are described in Section 4 and the Internet

Table 2. Sample construction by decade

	1992–1999	2000–2009	2010–2019	2020–2026
Raw bar rows	626,407,882	2,331,426,585	2,657,373,989	2,071,067,068
Duplicate rows merged	0	0	0	10,128
OHLC-violation flags	90	10	0	0
Split-window flagged rows	772,707	1,811,211	815,630	897,786
Regular-hours bar rows	611,399,568	2,278,758,476	2,599,857,420	1,940,283,624
Distinct symbols	26,519	22,052	18,761	46,985
Trading days	2,022	2,515	2,516	1,569

Notes: Full panel (all instruments). Duplicates merged by the rule open=first, high=max, low=min, close=last, volume=sum. Flags mark rows; no raw data are deleted. Split windows are ± 1 trading day around high-confidence overnight-ratio events (Section 3).




Figure 3. Coverage of the full panel. (a) Distinct symbols per year and the screened (lagged top-3000-eligible) universe; (b) regular-hours bars per year; (c) median and 10th-percentile bars per stock-day on the lagged top-1000 universe.

Appendix. In the pilot months the standard deviation of unscreened minute returns is an order of magnitude larger than the 1st–99th percentile spread, reflecting split-like jumps and isolated bad prints that the cleaning stage must remove.

4. Measurement

This section defines every object used in the empirical analysis, in the order in which the pipeline constructs them: returns, residuals, standardized innovations, source shocks, isolation filters, propagation outcomes, networks, and systemic co-jump minutes. Two principles govern all definitions. First, *strict lag discipline*: any parameter entering the transformation of month- m data—factor loadings, scale estimates, seasonal profiles, universe membership, network weights—is estimated exclusively on data dated before month m . Second, *measurement before interpretation*: each con-

struct is defined so that its mechanical properties (denominators, censoring, base rates) are known before any economic claim rests on it.

4.1. From bars to returns

The unit of observation is a one-minute bar during the regular session (09:30–16:00 ET, truncated at the true exchange close on early-closing days). Let $P_{i,t}$ be the last trade price of security i in minute t . One-minute log returns,

$$r_{i,t} = \log P_{i,t} - \log P_{i,t-1}, \quad (1)$$

are computed within ticker and trading day only: the first observed minute of each stock-day has no return, and overnight intervals are never bridged, so no intraday statistic mixes overnight information. Because the source records a bar only when at least one trade occurs, a missing minute means “no trade,” not “zero return”; returns are accordingly defined only across consecutively observed minutes and missingness is carried, never imputed. Bars flagged for OHLC inconsistency and all bars inside detected corporate-action windows (Section 3) are excluded before any return is computed.

All shock-level analysis operates on the lagged top-500 dollar-volume universe. This choice trades breadth for interpretability: top-500 names trade essentially every minute (median 382–390 bars per day from the mid-1990s onward), so the cross-sectional response measures defined below have stable, nearly complete denominators, and the asynchronicity that depresses high-frequency comovement among thinner names (Epps, 1979) cannot masquerade as absent propagation. Results for top-1000 and top-3000 universes appear in the robustness materials.

4.2. Residualization

A raw extreme return is not evidence of an idiosyncratic event: it may load on a market-wide movement. We therefore work with residuals from a single-factor market model whose ingredients are constructed to avoid three mechanical contaminations.

Leave-one-out factor. The market return at minute t , from the perspective of security i , is the equal-weighted mean return of the universe *excluding* i :

$$r_{m,t}^{(-i)} = \frac{1}{N_t - 1} \sum_{j \neq i} r_{j,t}, \quad (2)$$

where N_t counts securities with an observable return at t . Without the leave-one-out correction, a large own return mechanically moves the factor and attenuates the measured shock; with 500 names the correction is small but we apply it exactly.

Lagged loadings. The loading β_i is the OLS slope of $r_{i,t}$ on the equal-weighted market return, estimated on all minutes of month $m-1$ in which the security has at least 300 valid return observations and the market aggregate at least 50 contributing names; securities with insufficient history receive $\beta_i = 1$, and estimates are winsorized to $[-3, 5]$. Residuals in month m are then

$$e_{i,t} = r_{i,t} - \beta_i r_{m,t}^{(-i)}. \quad (3)$$

A single lagged loading per month is deliberately coarse: it cannot overfit the event minutes it will later be used to classify, and any misspecification it induces is absorbed by the placebo differencing described in Section 5.

Robust, seasonally adjusted scale. Minute returns are strongly leptokurtic and their dispersion follows the familiar intraday U-shape (Andersen and Bollerslev, 1997); a sample standard deviation would be dominated by the very tail events we wish to detect, and an unadjusted scale would over-detect shocks at the open. We therefore standardize by a median-absolute-deviation scale interacted with a minute-of-day factor, both from month $m-1$:

$$z_{i,t} = \frac{e_{i,t}}{1.4826 \text{MAD}_i \times f_{\tau(t)}}, \quad f_{\tau} = \frac{\text{med}_i |e_{i,\cdot}| / (1.4826 \text{MAD}_i) |_{\tau}}{\overline{\text{med}}}, \quad (4)$$

where $\tau(t)$ is the minute of day, the numerator of f_{τ} is the cross-sectional median absolute standardized residual at minute-of-day τ in month $m-1$, and $\overline{\text{med}}$ normalizes the profile to unit mean. The factor 1.4826 makes the MAD consistent for the standard deviation under normality; under the actual fat-tailed distribution, z is a ranking device whose thresholds are calibrated empirically, not a Gaussian quantile. The same construction applied to the market return, with the prior month's market MAD, yields the standardized market factor $z_{m,t}$ used in the isolation filters.

4.3. Source shocks and isolation

A *source shock* is an event (i, t) with

$$|z_{i,t}| \geq 6, \quad (5)$$

with the sign recorded and analyzed separately. The threshold is deliberately extreme: it selects on the order of one event per stock-week (1,854,384 events over 410 months), deep enough into the tail that bid-ask bounce alone cannot generate qualifying events for top-500 names after 2001, while leaving enough events for precise inference within every year. Nested thresholds ($|z| \geq 8$, $|z| \geq 12$) are reported in the robustness table.

Because the question concerns *idiosyncratic* origins, two isolation filters apply at the event minute:

1. **Calm market:** $|z_{m,t}| < 1$. The event does not coincide with a market-factor movement.
2. **No simultaneous wave:** the share of eligible securities with $|z_{j,t}| > 4$ at minute t is at most 2%, and at least 50 securities are observable. The event is not part of a market-wide burst.

Both filters are measured, not assumed: minutes excluded by filter (2) form the systemic co-jump sample of Section 4.6, so the design partitions all extreme activity into an idiosyncratic branch and a systemic branch rather than discarding the inconvenient one. The filters are intentionally conservative about common information; the cost, acknowledged throughout, is that any propagation operating *through* an immediate market-wide reaction is assigned to the systemic branch by construction.

4.4. Propagation outcomes

Fix an event (i, t) and a horizon $h \in \{1, 2, 5, 10, 15, 30, 60, 120\}$ minutes. Receivers are all $j \neq i$ in the universe with an observable bar at the measurement minute; outcomes require at least 50 eligible receivers and are right-censored—recorded as missing, not zero—whenever $t + h$ crosses the session close, so end-of-day events never contribute mechanically attenuated responses. The outcomes are:

$$\text{Breadth}_{i,t}(h) = \frac{1}{N_{t,h} - 1} \sum_{j \neq i} \mathbf{1}\{|z_{j,t+h}| > 2\}, \quad (6)$$

$$\text{Intensity}_{i,t}(h) = \frac{1}{N_{t,h} - 1} \sum_{j \neq i} |z_{j,t+h}|, \quad (7)$$

$$\text{Signed}_{i,t}(h) = \frac{\text{sign}(e_{i,t})}{N_{t,h} - 1} \sum_{j \neq i} z_{j,t+h}. \quad (8)$$

The receiver threshold $|z| > 2$ is calibrated to the fat-tailed base rate: it fires for roughly 7% of stock-minutes unconditionally, high enough that breadth is a well-populated share whose *changes* are measured with precision, low enough to respond to genuine tail activity. Cumulative breadth replaces the point-in-time indicator with $\mathbf{1}\{\max_{s \in (t, t+h]} |z_{j,s}| > 2\}$, so a receiver that jumps at $t + 3$ and is quiet at $t + 15$ still counts; its base rate is mechanically higher (about 53% at $h = 15$) and its pre-event benchmark uses the mirrored window $[t - 5 - h, t - 5]$.

Every outcome is also evaluated at $t - 5$ (point-in-time) or on the mirrored pre-window (cumulative). The pre-event value serves two distinct purposes: it is the *differencing baseline*—shocks arrive disproportionately in minutes of elevated ambient activity, and the post-minus-pre difference removes that level—and it is a *falsification statistic*, since genuine propagation cannot precede its cause.

4.5. Lagged correlation networks

For each month m and security i , the neighbor set $\mathcal{N}_i$ collects the twenty securities with the highest absolute pairwise correlation with i 's residuals over month $m-1$ (pairwise complete observations; neighbor sets require at least ten members present in month m to be used). Splitting equation (6) into

$$\text{Breadth}_{i,t}^{\mathcal{N}}(h) = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \mathbf{1}\{|z_{j,t+h}| > 2\}, \quad \text{Breadth}_{i,t}^{\overline{\mathcal{N}}}(h) = \text{same over } j \notin \mathcal{N}_i \cup \{i\}, \quad (9)$$

yields the *network differential* $\text{Breadth}^{\mathcal{N}} - \text{Breadth}^{\overline{\mathcal{N}}}$, the paper's sharpest instrument: transmission concentrated among twenty connected names moves it one-for-one, while market-wide breadth dilutes the same response by a factor of twenty-five. Because neighbor sets are lagged, the differential cannot reflect the event's own contemporaneous comovement; and because the identical differential is computed for placebo events, any unconditional tendency of correlated stocks to co-move nets out.

4.6. Systemic co-jump minutes

A minute is a *systemic co-jump* if the share of eligible securities with $|z_{j,t}| > 4$ exceeds 2% (with at least 50 eligible names)—exactly the complement of isolation filter (2). For each such minute we record the affected shares at the 4σ and 6σ thresholds and the standardized market return. Two summary series are used: the frequency of *severe* synchronization (affected share above 10%) per 1,000 eligible minutes, and the mean affected share within co-jump minutes. We emphasize severity and the severe-minute count rather than the raw count at the 2% threshold because the latter is not comparable across regimes: idiosyncratic tail activity was two orders of magnitude more frequent in the wide-tick 1990s, so a fixed low threshold is breached almost continuously early in the sample for reasons unrelated to synchronization.

4.7. What these measures can and cannot see

Three mechanical properties discipline interpretation. First, breadth measures respond to *tail* activity only; diffuse sub-threshold responses would appear in intensity, which is reported alongside breadth throughout. Second, all receiver measures condition on trading: a receiver without a bar at $t+h$ leaves the denominator rather than entering as a zero, so illiquidity cannot fabricate absorption. Third, the pre-2001 tick regime (eighths and sixteenths) inflates both the level of $|z|$ activity and its base rates; every cross-era comparison is therefore anchored on the within-era placebo contrast, and the post-decimalization subsample reproduces all headline results.

5. Empirical design

5.1. The estimand

For an outcome Y from Section 4.4, define the within-event change $\Delta Y_{i,t}(h) = Y_{i,t}(t+h) - Y_{i,t}(t-5)$. The estimand is the difference-in-differences

$$\Delta\Delta Y(h) = \mathbb{E}[\Delta Y(h) \mid \text{real event}] - \mathbb{E}[\Delta Y(h) \mid \text{matched placebo}]. \quad (10)$$

In potential-outcomes language, the placebo arm estimates the counterfactual path of cross-sectional activity that the event window would have followed had the source stock not jumped, holding fixed the stock, the minute of day, the month, and the calm-market state. Three biases that a naive event study would absorb are differenced out explicitly. *Volatility clustering*: shocks arrive in already-active minutes, so post-event activity is elevated even absent transmission; the inner difference removes the level. *Deterministic intraday drift*: breadth trends over the trading day; the placebo, matched on minute of day, follows the same drift. *Filter-induced selection*: conditioning on a calm market at t selects minutes followed by reversion toward normal activity; the placebo is drawn under the identical filters and inherits the identical reversion.

5.2. Placebo construction

For each real event (i, t) the matched placebo is drawn by the following algorithm, with a fixed seed recorded in the replication package:

1. collect all minutes of the same month with the same minute of day $\tau(t)$, on a different trading day, at which (a) security i has an observable return with $|z_{i,\tau}| < 2$, (b) the market-calm filter holds, and (c) the simultaneity filter holds;
2. draw one candidate uniformly at random; compute every outcome of Section 4.4 at the placebo minute by the identical code path;
3. label the placebo with the real event's magnitude and sign so that all conditional cuts (by $|z|$ bin, by sign, by state) compare real and placebo events within the same cell.

Events with an empty candidate set (rare) contribute to the real arm only and are dropped from difference-in-differences estimation by the pairing on dates within cells. The match rate exceeds 95% throughout.

5.3. Inference

Outcomes are correlated across events within a trading day—mechanically so, because events of the same minute share receivers. All inference therefore resamples *days*: for $b = 1, \dots, 500$, draw

event dates with replacement, recompute date-level sums and counts for both arms, and form the bootstrap distribution of the full statistic (10). Standard errors are bootstrap standard deviations; the month-cluster variant (resampling year-months) appears in the robustness table and roughly doubles the standard error of the network differential without changing any conclusion. With roughly 8,500 event days, the effective sample is large, and the design’s minimum detectable effect for market-wide breadth at $h = 5$ is approximately 2×10^{-4} at the 5% level with 80% power—one additional responding security per twenty-five events. Null results below are powered nulls.

5.4. Identification: what is and is not claimed

The design identifies the *association* between an isolated idiosyncratic jump and subsequent cross-sectional tail activity, purged of ambient levels, intraday seasonality, calm-market selection, and the stock’s own base rates. It does not identify the structural channel behind the network differential: lagged-correlation neighbors are typically industry peers, ETF co-members, and factor-exposure twins, and the differential aggregates fundamental information transmission (investors rationally repricing linked firms) with trading-channel transmission (basket flows, hedging, liquidity withdrawal). Separating those requires timestamped news and order-flow data and is beyond OHLCV bars; the paper’s claims are correspondingly phrased as predictive propagation, not causal transmission. Likewise, the systemic co-jump series conflates common news arrivals with endogenous synchronization; we document its evolution and severity without attributing minutes to either source.

5.5. Falsification battery

Five falsifiable predictions are tested and reported. (1) *No pre-event response*: real-event outcomes at $t - 5$ equal placebo levels. (2) *No placebo network structure*: the neighbor–non-neighbor differential computed on placebo events alone is zero ($-2.7 \pm 2.8 \times 10^{-4}$). (3) *No announcement artifact*: excluding the 10:00 and 14:00–14:30 ET windows that host most scheduled releases leaves all estimates unchanged. (4) *No coverage artifact*: fixed-rule lagged universes make the 2020 source-coverage expansion irrelevant to any comparison. (5) *No tick-regime artifact*: the post-decimalization subsample reproduces the market-wide null and the network differential. All five pass; Table 5 quantifies each.

6. Results

6.1. The transformation of the intraday market, 1992–2026

Figures 4 through 6 document the environment in which propagation must be measured, on liquidity-screened universes with strictly lagged membership.

Three secular changes stand out. First, ordinary minute-level variability collapsed. Mean absolute one-minute returns on the top-1000 universe fell decade by decade at every minute of the trading day (Figure 4a), and the frequency of one-minute returns exceeding 1% fell from about 518 per 10,000 in the 1990s to 4.8 per 10,000 in the 2010s—a factor of roughly 108 (Figure 5). Part of the 1990s level reflects bid–ask bounce under eighth and sixteenth ticks rather than true volatility; the decade-by-decade decline survives range-based measures and coarser sampling, and we flag the tick-regime caveat wherever it binds.

Second, the collapse was not monotone in the tails. The frequency of $|r| > 5\%$ minutes rose from 0.015 per 10,000 in 2017–2019 to 5.8 in 2020—a factor of 391—and 2008 shows the same signature (Figure 5). Extremes became rarer but, when the market is stressed, arrive in synchronized bursts; we return to this in Section 7.

Third, cross-sectional information transmission accelerated dramatically. In 1996 the mean pairwise correlation of one-minute returns among the hundred largest stocks was 0.004, against 0.16 at thirty minutes: comovement existed but took tens of minutes to materialize, the classical Epps effect (Epps, 1979). By 2015 the one-minute correlation was 0.30 against 0.39 at thirty minutes, and 2024 shows the same near-flat profile (Figure 6). The Epps effect has essentially disappeared from large-cap U.S. equities: common information is now impounded across the cross-section within a minute.

6.2. Baseline: idiosyncratic shocks do not propagate

Table 3 reports the paper’s baseline estimates of equation (10) for 1,854,384 isolated $|z| \geq 6$ source events and 1,766,058 matched placebos. The abnormal breadth response is economically zero at every horizon: at $h = 5$ minutes the point estimate is -1.04×10^{-4} (bootstrap SE 0.90×10^{-4}), and at $h = 30$ minutes -3.10×10^{-4} (SE 1.01×10^{-4}). Against a placebo base rate of roughly 7×10^{-2} , the largest estimate in absolute value implies that an isolated six-sigma shock changes the probability that any given other security jumps by less than half of one percent *of the base rate*—with the wrong sign. Abnormal intensity estimates are equally null.

Three features of Table 3 deserve emphasis. First, this is a powered null: standard errors are small enough to detect a breadth response of 3×10^{-4} —one additional responding security per seventeen events—at conventional levels. Second, the pre-event falsification passes: real-event




Figure 4. Intraday seasonality by decade. Minute-of-day means during regular hours, lagged top-1000 dollar-volume universe, split-window and flagged bars excluded: (a) absolute one-minute log return; (b) share of the stock-day’s volume; (c) high–low log range.

Table 3. Baseline abnormal propagation after idiosyncratic shocks

Horizon (min)	1	2	5	10	15	30	60	120
$\Delta\Delta$ Breadth ($\times 10^{-4}$)	-1.31 (0.85)	-0.85 (0.92)	-1.04 (0.90)	-0.27 (1.00)	-1.46 (0.98)	-3.10 (1.01)	-1.24 (1.21)	-2.69 (1.47)
$\Delta\Delta$ Intensity ($\times 10^{-2}$)	1.02 (1.43)	-0.66 (1.15)	1.43 (1.44)	1.61 (0.97)	0.23 (1.28)	-0.18 (1.37)	5.70 (2.35)	0.25 (1.34)
Events with valid h	1,827,839	1,822,057	1,806,236	1,782,163	1,757,825	1,687,295	1,545,783	1,246,819

Notes: $\Delta\Delta = [\text{outcome}(t+h) - \text{outcome}(t-5)]$ for real shocks minus the same difference for matched placebos (same stock, same minute of day, calm day). Breadth = share of eligible receivers with $|z| > 2$; intensity = mean $|z|$. Source shocks: $|z| \geq 6$, market-calm and simultaneity isolation filters, lagged top-500 universe, 1992–2026. Day-block bootstrap standard errors (500 replications) in parentheses.

outcomes at $t - 5$ equal placebo levels, so the null is not produced by anticipation contaminating the baseline. Third, the null is uniform: Figure 7 shows the same flat response profile in each of the four decades, so absorption is not a recent acquisition of electronic markets—though the composition of what is being absorbed changed radically, as the descriptive results show.

Conditioning does not resurrect average propagation (Figure 8): estimates by shock sign and by calm versus stressed market states are all within two standard errors of zero. The only suggestive pattern is monotonicity in magnitude—the largest shocks (12σ – 20σ) carry a positive but insignificant point estimate of 3.47×10^{-4} (SE 2.95×10^{-4})—which motivates the concentrated-propagation probes of the next section.

7. Mechanisms

The baseline null says that isolated idiosyncratic shocks do not raise market-wide jump activity. This section asks where transmission lives if not there: in the shocked stock’s network neighborhood,




Figure 5. Extreme one-minute returns. Annual frequency (log scale) of one-minute absolute log returns exceeding 1%, 2%, and 5%, lagged top-1000 universe, regular hours. Pre-2001 levels partly reflect eighth/sixteenth tick bid–ask bounce.

in transient responses between measurement points, and in the systemic co-jump minutes that the idiosyncratic design deliberately sets aside.

7.1. Network-concentrated propagation — and its intensification

Market-wide breadth spreads any response across five hundred receivers; a transmission channel operating through economic or statistical linkages concentrates it among a handful. For every event we therefore compare the response of the source’s twenty closest lagged-correlation neighbors with the response of all other receivers, real minus placebo (Table 4).

The differential is large, precisely estimated, and horizon-decaying: five minutes after an isolated $|z| \geq 6$ shock, the source’s neighbors are 64.0×10^{-4} more likely than non-neighbors to exceed $|z| > 2$ —relative to placebo—with a bootstrap standard error of 1.8×10^{-4} ($t \approx 34$); the differential declines to 49.6×10^{-4} at thirty minutes. Against the $\sim 700 \times 10^{-4}$ base rate, a neighbor’s jump probability rises by roughly nine percent of its unconditional level. Idiosyncratic shocks do propagate—to the connected few, not to the diffuse many, which is precisely why market-wide breadth cannot detect them.

The evolution across decades is the paper’s sharpest dynamic fact (Figure 9): the five-minute neighbor differential was 28×10^{-4} in 1992–1999 and 30×10^{-4} in 2000–2009, then jumps to 81×10^{-4} in 2010–2019 and 90×10^{-4} in 2020–2026—a tripling concentrated exactly in the era of algorithmic liquidity provision, correlation-aware execution, and index intermediation. Propagation along the statistical network became roughly three times stronger while market-wide spillover




Figure 6. The disappearance of the Epps effect. Mean pairwise return correlation among the top-100 lagged dollar-volume stocks at 1-, 5-, 15-, and 30-minute sampling, previous-tick synchronization within day, for 1996, 2005, 2015, and 2024.

Table 4. Network mechanism: neighbor minus non-neighbor abnormal response

Sample	$h=5$	$h=15$	$h=30$
Full sample	64.02 (1.82)	55.36 (1.80)	49.64 (1.85)
1992–1999	28.35 (7.16)	22.48 (6.96)	29.91 (7.38)
2000–2009	29.95 (2.23)	24.89 (2.11)	19.45 (2.44)
2010–2019	80.64 (3.47)	71.05 (3.30)	64.28 (3.51)
2020–2026	89.60 (4.37)	76.56 (4.40)	69.29 (4.56)

Notes: $\Delta\Delta$ of [breadth among the source’s top-20 lagged-correlation neighbors – breadth among non-neighbors] at $t+h$, real minus matched placebo, $\times 10^{-4}$. Day-block bootstrap SEs (500 reps) in parentheses. Neighbors from month $m-1$ residual correlations; source shocks $|z| \geq 6$ with isolation filters, top-500.

remained zero: transmission grew *sharper*, not broader.

7.2. Cumulative responses

Point-in-time breadth at $t+h$ misses a receiver that jumps at $t+3$ and is quiet at $t+15$. Cumulative breadth over $(t, t+h]$, differenced against its mirrored pre-window and the placebo, remains null at the market level (at thirty minutes: 0.7×10^{-4} , SE 1.9×10^{-4}). The market-wide null is therefore not an artifact of sampling the response at isolated instants; transient market-wide cascades after isolated idiosyncratic shocks are absent as well.




Figure 7. Abnormal propagation breadth by decade. Point estimates of $\Delta\Delta$ breadth (equation (10)) with 95% day-block bootstrap bands, by horizon and decade.




Figure 8. Shock size, sign, and market state. $\Delta\Delta$ breadth at $h = 15$ by source-shock magnitude bins, shock sign, and prior-30-minute market volatility state, with 95% bootstrap intervals.

7.3. Systemic co-jump minutes: where fragility lives

The isolation filters exclude minutes in which more than 2% of the universe jumps simultaneously. These minutes are the natural home of the fragility hypothesis, and their time series (Figure 10) delivers two messages. First, the ordinary cross-section *de-synchronized* secularly: with idiosyncratic tail activity a hundred times rarer, minutes breaching the fixed 2% threshold fell from near-ubiquity in the sparse, wide-tick 1990s to about one minute in eight today—a comparability warning we take seriously, and the reason the fixed-threshold count is not itself a fragility measure. Second, conditional severity moves the other way in stress: within co-jump minutes, the mean share of names beyond 4σ averages 3.5% in calm recent years but reaches 9.3% in 2008 and 6.4% in 2020, and severe-synchronization minutes (more than 10% of the universe jumping together) spike from




Figure 9. Network-concentrated propagation by decade. Neighbor minus non-neighbor $\Delta\Delta$ breadth at $h \in \{5, 15, 30\}$ with 95% day-block bootstrap intervals. Neighbors are the source’s top-20 absolute-correlation partners estimated on prior-month residuals.

9.53 per 1,000 in 2017–2019 to 83.7 per 1,000 in 2020. Combined with the figE5 evidence that $|r| > 5\%$ minutes rose by a factor of 391 in 2020, the pattern is consistent across measures: modern stress arrives as market-wide synchronized minutes, not as single-name events.

Taken together, the mechanism evidence resolves the baseline null into a sharper statement: idiosyncratic shocks propagate narrowly along the lagged correlation network, that channel strengthened threefold in the algorithmic era, and market-wide fragility is the province of common-shock minutes whose severity—not frequency—is the binding object.

8. Robustness

The Internet Appendix reports the full battery specified in the analysis protocol. The main text highlights the checks that discipline interpretation. *Tick-regime*: pre-2001 minute returns embed eighth/sixteenth-tick bid–ask bounce; all secular comparisons are re-run on range-based measures, on 5- and 15-minute aggregation, and within post-decimalization subsamples. *Coverage*: the 2020 source-coverage expansion is absorbed by fixed-rule lagged top- N universes; results are reported for $N \in \{500, 1000, 3000\}$. *Universe composition*: ETFs and likely non-common instruments (ticker-shape heuristic) are excluded in robustness runs. *Corporate actions*: split windows are excluded throughout; sensitivity to window width and detection thresholds is tabulated. *Thresholds*: source shocks at $|z| \geq 4, 5, 6$ and receiver thresholds at $|z| > 2, 3$ leave the baseline null unchanged. *Inference*: day-block and month-block bootstrap agree. Table 5 collects the headline checks: the market-wide null and the $+64 \times 10^{-4}$ network differential are stable across the post-decimalization subsample, mid-day-only events, announcement-window exclusions, calm and crisis years, higher




Figure 10. Systemic co-jump minutes, 1992–2026. (a) Minutes per 1,000 in which more than 10% of eligible top-500 names exceed $|z| > 4$ simultaneously (log scale). (b) Mean affected share within co-jump minutes ($> 2\%$ threshold).

source thresholds, and non-common-ticker exclusions, and the neighbor differential measured on placebo events alone is zero ($-2.7 \pm 2.8 \times 10^{-4}$)—the permuted-benchmark falsification the design commits to. The one systematic pattern is that the network differential *declines* with source magnitude (+64 at $|z| \geq 6$ to $+34 \times 10^{-4}$ at $|z| \geq 12$), consistent with the largest events being firm-specific news that neighbors correctly do not price.

9. Conclusion

Three decades of minute-level data deliver a sharper statement of intraday market fragility than either the efficiency or the fragility narrative alone. Isolated idiosyncratic shocks—even six-sigma events in the largest five hundred securities—do not propagate: their abnormal cross-sectional footprint is statistically and economically zero at every horizon from one minute to two hours, in every decade. What changed is the market’s handling of common variation: information that in the mid-1990s took half an hour to appear in cross-stock comovement now arrives within the minute, ordinary extreme moves became an order of magnitude rarer, and extreme synchronization concentrated in systemic co-jump minutes whose frequency explodes in stress episodes. Fragility migrated from securities to minutes.

The evidence is deliberately associational where identification runs out: OHLCV bars do not observe quotes, depth, or order flow, so liquidity operates through proxies; tickers are not permanent identities, so long-run statements run within screened spells; and the co-jump series conflates common news with endogenous synchronization—separating the two requires timestamped news and quote data, which we leave as the natural extension. For market design and risk monitoring, the

Table 5. Robustness of the market-wide null and the network channel

Sample / design	$\Delta\Delta\text{Breadth}_5$	$\Delta\Delta\text{Breadth}_{30}$	Network diff ₅
Full sample	-1.04 (0.87)	-3.10 (0.99)	64.02 (1.84)
Post-decimalization (2002+)	-1.03 (0.97)	-3.04 (1.25)	65.81 (1.96)
Mid-day events (10:00–15:30)	-1.12 (1.04)	-3.48 (1.07)	64.75 (2.03)
Ex announcement windows	-1.28 (0.96)	-3.50 (1.11)	63.34 (1.92)
Calm years (ex 2008–09, 2020)	-1.35 (0.87)	-3.29 (1.00)	63.62 (1.95)
Crisis years only	2.69 (5.14)	-0.79 (6.32)	68.06 (6.34)
Source $ z \geq 8$	0.61 (1.23)	-1.66 (1.43)	49.89 (2.09)
Source $ z \geq 12$	1.12 (1.45)	0.32 (1.78)	33.94 (2.28)
Ex non-common tickers	-1.04 (0.95)	-3.08 (1.06)	63.99 (1.87)
Month-cluster bootstrap	-1.04 (0.92)	—	64.02 (3.44)
Placebo events, neighbor diff	—	—	-2.71 (2.82)

Notes: All entries $\times 10^{-4}$, day-block bootstrap SEs in parentheses (500 reps; final row of the market-wide block uses month-cluster resampling). $\Delta\Delta\text{Breadth}$ = abnormal market-wide breadth (real – placebo, post – pre). Network diff = neighbor – non-neighbor abnormal breadth at $h=5$. ‘Placebo events’ row reports the neighbor differential among placebo events alone (falsification; expected ≈ 0).

practical implication is direct: surveillance aimed at single-name events monitors the part of the system that modern markets already absorb; the binding constraint is the market-wide minute.

References

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 2012, 80 (5), 1977–2016.
- Acharya, Viral V. and Lasse Heje Pedersen**, “Asset pricing with liquidity risk,” *Journal of Financial Economics*, 2005, 77 (2), 375–410.
- Aït-Sahalia, Yacine and Dacheng Xiu**, “Principal Component Analysis of High-Frequency Data,” *Journal of the American Statistical Association*, 2019, 114 (525), 287–303.
- **and Jean Jacod**, “Testing for Jumps in a Discretely Observed Process,” *Annals of Statistics*, 2009, 37 (1), 184–222.
- **, Julio Cacho-Diaz, and Roger J. A. Laeven**, “Modeling Financial Contagion Using Mutually Exciting Jump Processes,” *Journal of Financial Economics*, 2015, 117 (3), 585–606.
- Amihud, Yakov**, “Illiquidity and stock returns: cross-section and time-series effects,” *Journal of Financial Markets*, 2002, 5 (1), 31–56.
- Andersen, Torben G. and Oleg Bondarenko**, “VPIN and the Flash Crash,” *Journal of Financial Markets*, 2014, 17, 1–46.
- **and Tim Bollerslev**, “Intraday Periodicity and Volatility Persistence in Financial Markets,” *Journal of Empirical Finance*, 1997, 4 (2–3), 115–158.
- **, —, Francis X. Diebold, and Clara Vega**, “Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange,” *American Economic Review*, 2003, 93 (1), 38–62.
- **, —, —, and —**, “Real-Time Price Discovery in Global Stock, Bond and Foreign Exchange Markets,” *Journal of International Economics*, 2007, 73 (2), 251–277.
- Andrade, Sandro C., Charles Chang, and Mark S. Seasholes**, “Trading Imbalances, Predictable Reversals, and Cross-Stock Price Pressure,” *Journal of Financial Economics*, 2008, 88 (2), 406–423.
- Antón, Miguel and Christopher Polk**, “Connected Stocks,” *Journal of Finance*, 2014, 69 (3), 1099–1127.
- Aquilina, Matteo, Eric Budish, and Peter O’Neill**, “Quantifying the High-Frequency Trading “Arms Race”,” *Quarterly Journal of Economics*, 2022, 137 (1), 493–564.

- Bacry, Emmanuel, Iacopo Mastromatteo, and Jean-François Muzy**, “Hawkes Processes in Finance,” *Market Microstructure and Liquidity*, 2015, 1 (1), 1550005.
- Bajgrowicz, Pierre, Olivier Scaillet, and Adrien Treccani**, “Jumps in High-Frequency Data: Spurious Detections, Dynamics, and News,” *Management Science*, 2016, 62 (8), 2198–2217.
- Baldauf, Markus and Joshua Mollner**, “High-Frequency Trading and Market Performance,” *Journal of Finance*, 2020, 75 (3), 1495–1526.
- Baltussen, Guido, Sjoerd van Bakkum, and Zhi Da**, “Indexing and stock market serial dependence around the world,” *Journal of Financial Economics*, 2019, 132 (1), 26–48.
- Barigozzi, Matteo and Christian Brownlees**, “NETS: Network Estimation for Time Series,” *Journal of Applied Econometrics*, 2019, 34 (3), 347–364.
- Barndorff-Nielsen, Ole E. and Neil Shephard**, “Power and Bipower Variation with Stochastic Volatility and Jumps,” *Journal of Financial Econometrics*, 2004, 2 (1), 1–37.
- Barrot, Jean-Noël and Julien Sauvagnat**, “Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks,” *Quarterly Journal of Economics*, 2016, 131 (3), 1543–1592.
- Bekaert, Geert, Michael Ehrmann, Marcel Fratzscher, and Arnaud Mehl**, “The Global Crisis and Equity Market Contagion,” *Journal of Finance*, 2014, 69 (6), 2597–2649.
- Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi**, “Do ETFs Increase Volatility?,” *Journal of Finance*, 2018, 73 (6), 2471–2535.
- Bessembinder, Hendrik**, “Trade Execution Costs and Market Quality after Decimalization,” *Journal of Financial and Quantitative Analysis*, 2003, 38 (4), 747–777.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas**, “Equilibrium fast trading,” *Journal of Financial Economics*, 2015, 116 (2), 292–313.
- Billio, Monica, Mila Getmansky, Andrew W. Lo, and Lorian Pelizzon**, “Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors,” *Journal of Financial Economics*, 2012, 104 (3), 535–559.
- Boehmer, Ekkehart, Kingsley Fong, and Juan Julie Wu**, “Algorithmic Trading and Market Quality: International Evidence,” *Journal of Financial and Quantitative Analysis*, 2021, 56 (8), 2659–2688.
- Bollerslev, Tim and Viktor Todorov**, “Tails, Fears, and Risk Premia,” *Journal of Finance*, 2011, 66 (6), 2165–2211.

- , **Jia Li, Andrew J. Patton, and Rogier Quaedvlieg**, “Realized Semicovariances,” *Econometrica*, 2020, 88 (4), 1515–1551.
- , **Sophia Zhengzi Li, and Viktor Todorov**, “Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns,” *Journal of Financial Economics*, 2016, 120 (3), 464–490.
- , **Tzuo Hann Law, and George Tauchen**, “Risk, Jumps, and Diversification,” *Journal of Econometrics*, 2008, 144 (1), 234–256.
- Boudt, Kris and Mikaël Petitjean**, “Intraday Liquidity Dynamics and News Releases around Price Jumps: Evidence from the DJIA Stocks,” *Journal of Financial Markets*, 2014, 17, 121–149.
- Brogaard, Jonathan, Allen Carrion, Thibaut Moyaert, Ryan Riordan, Andriy Shkilko, and Konstantin Sokolov**, “High Frequency Trading and Extreme Price Movements,” *Journal of Financial Economics*, 2018, 128 (2), 253–265.
- , **Terrence Hendershott, and Ryan Riordan**, “High-Frequency Trading and Price Discovery,” *Review of Financial Studies*, 2014, 27 (8), 2267–2306.
- Brunnermeier, Markus K. and Lasse Heje Pedersen**, “Market Liquidity and Funding Liquidity,” *Review of Financial Studies*, 2009, 22 (6), 2201–2238.
- Budish, Eric, Peter Cramton, and John Shim**, “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response,” *Quarterly Journal of Economics*, 2015, 130 (4), 1547–1621.
- Caporin, Massimiliano, Aleksey Kolokolov, and Roberto Renò**, “Systemic Co-jumps,” *Journal of Financial Economics*, 2017, 126 (3), 563–591.
- Cespa, Giovanni and Thierry Foucault**, “Illiquidity Contagion and Liquidity Crashes,” *Review of Financial Studies*, 2014, 27 (6), 1615–1660.
- Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega**, “Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market,” *Journal of Finance*, 2014, 69 (5), 2045–2084.
- Chinco, Alex, Adam D. Clark-Joseph, and Mao Ye**, “Sparse Signals in the Cross-Section of Returns,” *Journal of Finance*, 2019, 74 (1), 449–492.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam**, “Commonality in liquidity,” *Journal of Financial Economics*, 2000, 56 (1), 3–28.

- , —, and —, “Liquidity and market efficiency,” *Journal of Financial Economics*, 2008, 87 (2), 249–268.
- , —, and —, “Recent trends in trading activity and market quality,” *Journal of Financial Economics*, 2011, 101 (2), 243–263.
- Christensen, Kim, Roel C. A. Oomen, and Mark Podolskij**, “Fact or Friction: Jumps at Ultra High Frequency,” *Journal of Financial Economics*, 2014, 114 (3), 576–599.
- Cohen, Lauren and Andrea Frazzini**, “Economic Links and Predictable Returns,” *Journal of Finance*, 2008, 63 (4), 1977–2011.
- Coles, Jeffrey L., Davidson Heath, and Matthew C. Ringgenberg**, “On index investing,” *Journal of Financial Economics*, 2022, 145 (3), 665–683.
- Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes**, “Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues,” *Journal of Finance*, 2010, 65 (1), 295–331.
- Corwin, Shane A. and Paul Schultz**, “A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices,” *Journal of Finance*, 2012, 67 (2), 719–760.
- Coughenour, Jay F. and Mohsen M. Saad**, “Common Market Makers and Commonality in Liquidity,” *Journal of Financial Economics*, 2004, 73 (1), 37–69.
- Coval, Joshua D. and Erik Stafford**, “Asset Fire Sales (and Purchases) in Equity Markets,” *Journal of Financial Economics*, 2007, 86 (2), 479–512.
- Da, Zhi and Sophie Shive**, “Exchange traded funds and asset return correlations,” *European Financial Management*, 2018, 24 (1), 136–168.
- Das, Sanjiv Ranjan and Raman Uppal**, “Systemic Risk and International Portfolio Choice,” *Journal of Finance*, 2004, 59 (6), 2809–2834.
- Diebold, Francis X. and Kamil Yilmaz**, “Better to Give than to Receive: Predictive Directional Measurement of Volatility Spillovers,” *International Journal of Forecasting*, 2012, 28 (1), 57–66.
- and —, “On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms,” *Journal of Econometrics*, 2014, 182 (1), 119–134.
- Duarte, Fernando and Thomas M. Eisenbach**, “Fire-Sale Spillovers and Systemic Risk,” *Journal of Finance*, 2021, 76 (3), 1251–1294.

- Easley, David, Marcos López de Prado, Maureen O'Hara, and Zhibai Zhang**, "Microstructure in the Machine Age," *Review of Financial Studies*, 2021, 34 (7), 3316–3363.
- , **Marcos M. López de Prado, and Maureen O'Hara**, "Flow Toxicity and Liquidity in a High-Frequency World," *Review of Financial Studies*, 2012, 25 (5), 1457–1493.
- Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson**, "Financial Networks and Contagion," *American Economic Review*, 2014, 104 (10), 3115–3153.
- Engle, Robert F., Takatoshi Ito, and Wen-Ling Lin**, "Meteor Showers or Heat Waves? Heteroskedastic Intra-Daily Volatility in the Foreign Exchange Market," *Econometrica*, 1990, 58 (3), 525–542.
- Epps, Thomas W.**, "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 1979, 74 (366), 291–298.
- Forbes, Kristin J. and Roberto Rigobon**, "No Contagion, Only Interdependence: Measuring Stock Market Comovements," *Journal of Finance*, 2002, 57 (5), 2223–2261.
- Foucault, Thierry, Johan Hombert, and Ioanid Roşu**, "News Trading and Speed," *Journal of Finance*, 2016, 71 (1), 335–382.
- , **Roman Kozhan, and Wing Wah Tham**, "Toxic Arbitrage," *Review of Financial Studies*, 2017, 30 (4), 1053–1094.
- Garman, Mark B. and Michael J. Klass**, "On the Estimation of Security Price Volatilities from Historical Data," *Journal of Business*, 1980, 53 (1), 67–78.
- Glasserman, Paul and H. Peyton Young**, "Contagion in Financial Networks," *Journal of Economic Literature*, 2016, 54 (3), 779–831.
- Glosten, Lawrence, Suresh Nallareddy, and Yuan Zou**, "ETF Activity and Informational Efficiency of Underlying Securities," *Management Science*, 2021, 67 (1), 22–47.
- Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka**, "Do liquidity measures measure liquidity?," *Journal of Financial Economics*, 2009, 92 (2), 153–181.
- Greenwood, Robin and David Thesmar**, "Stock Price Fragility," *Journal of Financial Economics*, 2011, 102 (3), 471–490.
- , **Augustin Landier, and David Thesmar**, "Vulnerable Banks," *Journal of Financial Economics*, 2015, 115 (3), 471–485.

- Gu, Shihao, Bryan Kelly, and Dacheng Xiu**, “Empirical Asset Pricing via Machine Learning,” *Review of Financial Studies*, 2020, 33 (5), 2223–2273.
- Hamao, Yasushi, Ronald W. Masulis, and Victor Ng**, “Correlations in Price Changes and Volatility across International Stock Markets,” *Review of Financial Studies*, 1990, 3 (2), 281–307.
- Hameed, Allaudeen, Wenjin Kang, and S. Viswanathan**, “Stock Market Declines and Liquidity,” *Journal of Finance*, 2010, 65 (1), 257–293.
- Hasbrouck, Joel**, “Measuring the Information Content of Stock Trades,” *Journal of Finance*, 1991, 46 (1), 179–207.
- **and Duane J. Seppi**, “Common factors in prices, order flows, and liquidity,” *Journal of Financial Economics*, 2001, 59 (3), 383–411.
- **and Gideon Saar**, “Low-Latency Trading,” *Journal of Financial Markets*, 2013, 16 (4), 646–679.
- Hautsch, Nikolaus, Julia Schaumburg, and Melanie Schienle**, “Financial Network Systemic Risk Contributions,” *Review of Finance*, 2015, 19 (2), 685–738.
- Hawkes, Alan G.**, “Spectra of Some Self-Exciting and Mutually Exciting Point Processes,” *Biometrika*, 1971, 58 (1), 83–90.
- Hendershott, Terrence and Pamela C. Moulton**, “Automation, speed, and stock market quality: The NYSE’s Hybrid,” *Journal of Financial Markets*, 2011, 14 (4), 568–604.
- **, Charles M. Jones, and Albert J. Menkveld**, “Does Algorithmic Trading Improve Liquidity?,” *Journal of Finance*, 2011, 66 (1), 1–33.
- Herskovic, Bernard**, “Networks in Production: Asset Pricing Implications,” *Journal of Finance*, 2018, 73 (4), 1785–1818.
- Hou, Kewei**, “Industry Information Diffusion and the Lead-Lag Effect in Stock Returns,” *Review of Financial Studies*, 2007, 20 (4), 1113–1138.
- Huth, Nicolas and Frédéric Abergel**, “High Frequency Lead/lag Relationships — Empirical Facts,” *Journal of Empirical Finance*, 2014, 26, 41–58.
- Israeli, Doron, Charles M. C. Lee, and Suhas A. Sridharan**, “Is there a dark side to exchange traded funds? An information perspective,” *Review of Accounting Studies*, 2017, 22 (3), 1048–1083.

- Jordà, Òscar**, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 2005, 95 (1), 161–182.
- Karolyi, G. Andrew, Kuan-Hui Lee, and Mathijs A. van Dijk**, “Understanding commonality in liquidity around the world,” *Journal of Financial Economics*, 2012, 105 (1), 82–112.
- King, Mervyn A. and Sushil Wadhwani**, “Transmission of Volatility between Stock Markets,” *Review of Financial Studies*, 1990, 3 (1), 5–33.
- Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun**, “The Flash Crash: High-Frequency Trading in an Electronic Market,” *Journal of Finance*, 2017, 72 (3), 967–998.
- Kodres, Laura E. and Matthew Pritsker**, “A Rational Expectations Model of Financial Contagion,” *Journal of Finance*, 2002, 57 (2), 769–799.
- Kyle, Albert S. and Wei Xiong**, “Contagion as a Wealth Effect,” *Journal of Finance*, 2001, 56 (4), 1401–1440.
- Lahaye, Jérôme, Sébastien Laurent, and Christopher J. Neely**, “Jumps, Cojumps and Macro Announcements,” *Journal of Applied Econometrics*, 2011, 26 (6), 893–921.
- Lee, Suzanne S. and Per A. Mykland**, “Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics,” *Review of Financial Studies*, 2008, 21 (6), 2535–2563.
- Lo, Andrew W. and A. Craig MacKinlay**, “When Are Contrarian Profits Due to Stock Market Overreaction?,” *Review of Financial Studies*, 1990, 3 (2), 175–205.
- Lou, Dong**, “A Flow-Based Explanation for Return Predictability,” *Review of Financial Studies*, 2012, 25 (12), 3457–3489.
- Lucca, David O. and Emanuel Moench**, “The Pre-FOMC Announcement Drift,” *Journal of Finance*, 2015, 70 (1), 329–371.
- Madhavan, Ananth**, “Exchange-Traded Funds, Market Structure, and the Flash Crash,” *Financial Analysts Journal*, 2012, 68 (4), 20–35.
- Menkveld, Albert J.**, “High Frequency Trading and the New Market Makers,” *Journal of Financial Markets*, 2013, 16 (4), 712–740.
- , “The Economics of High-Frequency Trading: Taking Stock,” *Annual Review of Financial Economics*, 2016, 8 (1), 1–24.

- **and Bart Zhou Yueshen**, “The Flash Crash: A Cautionary Tale About Highly Fragmented Markets,” *Management Science*, 2019, 65 (10), 4470–4488.
- Menzly, Lior and Oguzhan Ozbas**, “Market Segmentation and Cross-Predictability of Returns,” *Journal of Finance*, 2010, 65 (4), 1555–1580.
- Nagel, Stefan**, “Evaporating Liquidity,” *Review of Financial Studies*, 2012, 25 (7), 2005–2039.
- O’Hara, Maureen**, “High Frequency Market Microstructure,” *Journal of Financial Economics*, 2015, 116 (2), 257–270.
- **and Mao Ye**, “Is market fragmentation harming market quality?,” *Journal of Financial Economics*, 2011, 100 (3), 459–474.
- Parkinson, Michael**, “The Extreme Value Method for Estimating the Variance of the Rate of Return,” *Journal of Business*, 1980, 53 (1), 61–65.
- Pástor, Ľuboš and Robert F. Stambaugh**, “Liquidity Risk and Expected Stock Returns,” *Journal of Political Economy*, 2003, 111 (3), 642–685.
- Pelger, Markus**, “Understanding Systematic Risk: A High-Frequency Approach,” *Journal of Finance*, 2020, 75 (4), 2179–2220.
- Roll, Richard**, “A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market,” *Journal of Finance*, 1984, 39 (4), 1127–1139.
- Savor, Pavel and Mungo Wilson**, “How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements,” *Journal of Financial and Quantitative Analysis*, 2013, 48 (2), 343–375.
- Sirignano, Justin and Rama Cont**, “Universal features of price formation in financial markets: perspectives from deep learning,” *Quantitative Finance*, 2019, 19 (9), 1449–1459.
- van Kervel, Vincent and Albert J. Menkveld**, “High-Frequency Trading around Large Institutional Orders,” *Journal of Finance*, 2019, 74 (3), 1091–1137.
- Weller, Brian M.**, “Does Algorithmic Trading Reduce Information Acquisition?,” *Review of Financial Studies*, 2018, 31 (6), 2184–2226.
- Yuan, Kathy**, “Asymmetric Price Movements and Borrowing Constraints: A Rational Expectations Equilibrium Model of Crises, Contagion, and Confusion,” *Journal of Finance*, 2005, 60 (1), 379–411.