



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

October 2026

Version 1.1

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 findings emerge. First, isolated idiosyncratic shocks barely spread: across 1.85 million $\pm 6\sigma$ one-minute events, the abnormal share of other securities with an extreme standardized return, measured against matched within-stock placebos, lies between zero and about 2% of its base rate 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 ($+64.0 \times 10^{-4}$ in the probability of an extreme move at five minutes), and this network channel roughly tripled between the 1990s and the 2020s while market-wide spillover stayed bounded near zero—a rise that partly reflects sharper identification of the network as the Epps effect receded. Third, common information is impounded far faster (one-minute pairwise correlation 0.004 in 1996 versus 0.30 in 2015), ordinary extreme minutes became an order of magnitude rarer after decimalization, and severe co-jump minutes—more than a tenth of the universe moving together—are roughly ten times more frequent in 2008 and 2020 than in calm years. Modern markets absorb the idiosyncratic, transmit it narrowly along statistical linkages, and synchronize on 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 within a trading minute, does the shock stay local or does it spread? The answer shapes 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: May 6, 2010 showed that modern equity markets can move from normal functioning to synchronized collapse within minutes, faster than any human intervention (Kirilenko et al., 2017; Menkveld and Yueshen, 2019). 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 changed over the three decades in which trading moved from exchange floors to fragmented, algorithmic, ETF-intermediated electronic markets is, to our knowledge, undocumented.

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 (Hendershott et al., 2011; Menkveld, 2013; Brogaard et al., 2014). 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 (Brogaard et al., 2018; Cespa and Foucault, 2014; Ben-David et al., 2018). Modern markets may therefore absorb ordinary shocks more efficiently while synchronizing more tightly on common ones. The title asks whether U.S. equity markets became faster, wider, and more fragile at the intraday horizon. The answer this paper documents is: faster, yes; wider, no—transmission became sharper rather than broader; more fragile, only in a specific and bounded sense that we make precise.

Our laboratory is a panel of 7,686,275,524 one-minute bars 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, identifies securities by ticker alone, and does not document its venue scope, 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 and every limitation we could not remove.

Three findings organize the paper. First, *isolated idiosyncratic shocks do not raise market-wide tail activity by more than a few percent of its base rate*. We detect 1.85 million one-minute events in

which a top-500 stock moves by at least six robust standard deviations of its own residual return while the market factor is calm and fewer than 2% of other names move simultaneously. Relative to matched within-stock placebos, the difference-in-differences estimate of the share of other securities with a standardized residual beyond two in absolute value is between -1 and -3×10^{-4} at horizons of 1 to 120 minutes, against a base rate of about 0.10, with day-block bootstrap standard errors near 1×10^{-4} . We also show, however, that real events occur in minutes of elevated ambient cross-sectional activity (pre-event breadth is 40.3×10^{-4} higher than for placebos), so that the difference-in-differences and a pre-level-stratified comparison bracket the market-wide response between zero and roughly 2% of the base rate. Both bounds are economically small, and the pattern holds in every decade and for positive and negative shocks alike.

Second, *what transmission exists is concentrated and, within the limits of our measurement, has intensified*. The shocked stock’s twenty closest lagged-correlation neighbors are 64.0×10^{-4} more likely than other receivers to show a tail exceedance five minutes later, net of the same contrast for placebos—a response of about six percent of the neighbor base rate, estimated with a standard error of 1.8×10^{-4} (4.7×10^{-4} with ticker clustering). The differential is 28×10^{-4} in the 1990s, 30 in the 2000s, 81 in the 2010s, and 90 in the 2020s. We present this increase with an explicit caveat: neighbor sets are identified from one-minute residual correlations in the preceding month, and the paper’s own Epps evidence shows that those correlations were close to zero in the 1990s, so part of the secular rise reflects sharper identification of the network rather than stronger transmission along it.

Third, *common information is impounded across the cross-section far faster than it was, and tail risk within the post-decimalization era is concentrated in market-wide minutes*. In 1996 the mean pairwise correlation of one-minute returns among the hundred largest stocks was 0.004 against 0.16 at thirty minutes—the classical Epps effect—whereas in 2015 it was 0.30 against 0.39. 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; part of that decline is the mechanical disappearance of eighth and sixteenth ticks, and the decline measured from 2002 onward is roughly tenfold. Minutes in which more than 10% of the universe moves beyond four robust standard deviations simultaneously—severe co-jump minutes—are roughly an order of magnitude more frequent in 2008 and 2020 than in surrounding calm years, and the frequency of $|r| > 5\%$ minute returns rose from 0.015 per 10,000 in 2017–2019 to 5.8 in 2020. We do not claim that such synchronization is secularly more frequent than in the 1990s: under the same standardization the wide-tick 1990s produce severe-minute counts that exceed 2008, which we attribute to discreteness rather than fragility and which precludes a cross-regime comparison.

Identification receives explicit attention throughout. The central threat is that securities respond to common information rather than to each other. The shock definition residualizes against a leave-

one-out market factor with strictly lagged loadings and scales; isolation filters require a calm market factor 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 level of every outcome is reported rather than assumed equal across arms. The minutes that the isolation filters remove are not discarded but studied separately as the systemic co-jump sample. What the design identifies is the predictive association between an isolated idiosyncratic extreme and subsequent cross-sectional tail activity; it does not separate fundamental information transmission from trading-channel transmission, and we phrase every claim accordingly.

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 a uniform time series of propagation and synchronization measured across the entire structural transformation of U.S. trading. To the co-jump literature, we contribute a long-sample account of how the frequency and severity of synchronized extreme minutes evolved, with the comparability limits of such a series made explicit. And to the measurement literature, we document what OHLCV bars can and cannot identify about transmission, with every limitation quantified.

Section 2 positions the paper. Section 3 describes and audits the data. Section 4 defines shocks, outcomes, and networks. Section 5 presents the empirical design and the inference. Section 6 reports the descriptive transformation of the market and the baseline propagation results, including the design diagnostics. Section 7 examines networks, cumulative responses, and co-jumps. Section 8 summarizes robustness, and Section 9 concludes.

2. Related literature

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 and liquidity commonality; lead-lag effects and the Epps effect; network approaches to contagion; the transformation of U.S. market structure and the high-frequency trading debate; and the econometrics of jumps, co-jumps, and scheduled announcements. 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 agents who infer information from other markets' price changes, so that an idiosyncratic "mistake" in one market transmits broadly,

and [Kodres and Pritsker \(2002\)](#) show that cross-market rebalancing of shared macro exposures transmits idiosyncratic shocks even without common fundamentals. The second family operates through balance sheets and funding. [Kyle and Xiong \(2001\)](#) generate contagion as a wealth effect of convergence traders; [Yuan \(2005\)](#) shows that borrowing constraints make contagion and asymmetric price movements likelier after losses; and [Brunnermeier and Pedersen \(2009\)](#) link market liquidity to funding liquidity in margin spirals that explain sudden, correlated liquidity dry-ups. Most directly related to our 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 to May 6, 2010. These theories jointly predict that transmission should be state-dependent, asymmetric, and increasing in the degree to which pricing and liquidity provision are cross-security activities. Our measures of propagation breadth, network concentration, sign asymmetry, and market-state dependence are designed to test 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: forced sales by mutual funds facing outflows generate price pressure and reversals ([Coval and Stafford, 2007](#)), predictable flows generate predictable pressure ([Lou, 2012](#)), and uninformative imbalances in one stock spill over to correlated stocks through risk-averse liquidity providers ([Andrade et al., 2008](#)). Ownership structure predicts fragility ([Greenwood and Thesmar, 2011](#)), common active ownership forecasts correlation beyond fundamentals ([Antón and Polk, 2014](#)), and overlapping portfolios aggregate into measurable systemic vulnerability ([Greenwood et al., 2015](#); [Duarte and Eisenbach, 2021](#)). This work establishes that idiosyncratic demand shocks do propagate, but at monthly or quarterly resolution. Whether the same propagation operates within the hour, and whether its speed changed as intermediation became algorithmic, is the margin our data address.

2.3. Comovement, commonality, and liquidity dynamics

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 ([Coughenour and Saad, 2004](#)). Liquidity falls after market declines, especially when funding is tight ([Hameed et al., 2010](#)), specialist losses forecast spread widening ([Comerton-Forde et al., 2010](#)), and returns to liquidity provision spike 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 markets ([Karolyi et al., 2012](#)). Relative to this literature we do not

measure liquidity levels: bars contain no quotes, and the only liquidity information we use is prior-month dollar volume as a ranking device, in the spirit of [Amihud \(2002\)](#), and the high–low range as a descriptive volatility proxy ([Parkinson, 1980](#); [Garman and Klass, 1980](#)). Such OHLC-based proxies have a long pedigree ([Roll, 1984](#); [Corwin and Schultz, 2012](#); [Goyenko et al., 2009](#)), but they do not identify the quote-level mechanisms that the commonality literature studies.

2.4. Lead–lag effects, asynchronous trading, and the Epps effect

Cross-autocorrelation in returns is among the oldest facts in empirical asset pricing: large stocks lead small stocks ([Lo and MacKinlay, 1990b](#)), predominantly within industries ([Hou, 2007](#)), along customer–supplier links ([Cohen and Frazzini, 2008](#); [Menzly and Ozbas, 2010](#)), and, at tick resolution, from liquid to illiquid assets with intraday-seasonal structure ([Huth and Abergel, 2014](#)); [Chinco et al. \(2019\)](#) show that one-minute-ahead cross-stock predictability is sparse and short-lived. At the shortest horizons measured comovement is attenuated mechanically. [Epps \(1979\)](#) first documented that correlations among stocks fall toward zero as the sampling interval shrinks. Part of this “Epps effect” is a pure artifact of nonsynchronous trading and price discreteness ([Scholes and Williams, 1977](#); [Lo and MacKinlay, 1990a](#); [Harris, 1991](#)), which [Renò \(2003\)](#) shows can reproduce the decay by simulation and which estimators built on actual transaction times remove ([Hayashi and Yoshida, 2005](#); [Zhang, 2011](#)); microstructure noise further contaminates high-frequency second moments ([Hansen and Lunde, 2006](#)). Part is genuine delay in the incorporation of common information. [Tóth and Kertész \(2009\)](#) document that the Epps curve among NYSE stocks flattened steadily between 1993 and 2003 and argue that denser trading explains only part of the change. Our Epps evidence extends this series to 2024 on previous-tick synchronized grids and must be read with both mechanisms in mind; we return to the point in [Section 6](#). The same mechanics matter for our design, because neighbor sets are identified from one-minute residual correlations and are therefore identified less precisely in eras where the Epps effect was strong.

2.5. Networks, connectedness, and contagion measurement

Network theory shows when idiosyncratic shocks fail to diversify away ([Acemoglu et al., 2012](#); [Elliott et al., 2014](#)); [Glasserman and Young \(2016\)](#) survey the contagion-network literature; real-economy links propagate firm-level shocks measurably ([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 ([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](#)) and defines

contagion as comovement beyond fundamentals (Bekaert et al., 2014); classic volatility-spillover evidence (Engle et al., 1990; Hamao et al., 1990) anticipates our design at coarser resolution and for a handful of markets. These networks are estimated at daily or lower frequency on small panels. We estimate pre-shock correlation neighborhoods on minute data for the five hundred most active securities each month and use them as lagged conditioning information, with the same neighborhood contrast computed on placebo events as the benchmark.

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

Our sample spans the modern transformation of U.S. equity trading. Decimalization in 2001 cut spreads (Bessembinder, 2003); liquidity promotes efficiency (Chordia et al., 2008); turnover exploded in the 2000s alongside algorithmic trading (Chordia et al., 2011); NYSE automation made prices more efficient while raising the cost of immediacy (Hendershott and Moulton, 2011); and fragmentation did not measurably harm quality (O’Hara and Ye, 2011). The normal-times evidence on algorithmic and high-frequency trading is broadly positive (Hendershott et al., 2011; Brogaard et al., 2014; Menkveld, 2013; Hasbrouck and Saar, 2013; Chaboud et al., 2014; Boehmer et al., 2021), with Menkveld (2016) surveying the channel-by-channel evidence, and markets absorb even large predictable trades with limited lasting impact (Bessembinder et al., 2016). Theory tempers the optimism: speed investment can be socially excessive (Biais et al., 2015), fast news trading concentrates volume and short-run 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; Aquilina et al., 2022), and speed can crowd out information acquisition (Weller, 2018; Baldauf and Mollner, 2020), with fast traders 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 but measures liquidity levels, price discovery, or single events. We contribute the missing time series: propagation breadth and network concentration of extreme shocks measured uniformly across three decades of structural change, with year-by-year estimates rather than imposed trend breaks.

2.7. Fragility and flash events

Episodic evidence shows that 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 (Menkveld and Yueshen, 2019); and fragmentation shaped which ETFs were hit (Madhavan, 2012). Sub-second “fractures” in individual stocks proliferated as machine ecologies moved beyond human response times (Johnson et al., 2013). More systematically, high-frequency traders 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 (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 benchmarked against matched placebo minutes.

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 (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 (Israeli et al., 2017; Glosten et al., 2021; Coles et al., 2022). For us, the ETF layer is a candidate mechanism for the sharpening of propagation along correlation neighborhoods, since basket arbitrage converts idiosyncratic shocks into correlated flows. We cannot test it directly: the data carry no security-type labels, so ETFs are neither identified nor excluded, a limitation we state rather than paper over.

2.9. Jumps, co-jumps, and scheduled announcements

Our shock measurement builds on high-frequency econometrics. Intraday volatility periodicity must be filtered before anything else (Andersen and Bollerslev, 1997), and robust periodicity estimators are needed because the periodicity estimate is itself contaminated by the jumps one wishes to detect (Boudt et al., 2011); our minute-of-day factor is a cross-sectional median for that reason. 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), and jump tails carry extreme cross-sectional dependence (Bollerslev et al., 2013). Two cautions discipline our thresholds: at ultra-high frequency much of what coarser methods call jumps is friction (Christensen et al., 2014) or a “drift burst” extending over several minutes (Christensen et al., 2022), and multiple-testing corrections eliminate a large share of spurious detections (Bajgrowicz et al., 2016). Our events are threshold exceedances of standardized residuals, not test-based jump detections, and we use the term “extreme move” rather than “jump” wherever the distinction matters.

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); formal tests for common jump arrivals exist (Jacod and Todorov, 2009); co-jumps among large U.S. stocks are frequent and concentrated around announcements (Gilder

et al., 2014); simultaneous multivariate jumps are 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). Closest to our co-jump series, Bormetti et al. (2015) model systemic price co-jumps among many stocks with Hawkes factor processes and Calcagnile et al. (2018) document that the number of stocks jumping in the same minute is far larger than independence or a single factor would imply, with collective synchronization events increasing in frequency in the high-frequency era. 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) and tested at high frequency by Dungey et al. (2018). Scheduled macro announcements provide the identified common shocks against which idiosyncratic propagation must be benchmarked: announcement surprises produce jump-like responses (Andersen et al., 2003, 2007), generate co-jumps (Lahaye et al., 2011), structure volume and volatility (Bollerslev et al., 2018), and structure expected returns around announcement times (Savor and Wilson, 2013; Lucca and Moench, 2015). Our robustness battery accordingly excludes the main announcement windows. Where the co-jump literature tests for simultaneity, our propagation measures trace the sequence—breadth at one through 120 minutes—and our co-jump series traces the secular evolution of simultaneity itself.

2.10. Factor structure at the minute scale

Separating idiosyncratic from systematic movement requires a factor structure: common factors exist in high-frequency returns and order flow (Hasbrouck and Seppi, 2001), high-frequency principal components are consistent with low-frequency covariance structure (Aït-Sahalia and Xiu, 2019), latent continuous and jump factors are estimable (Pelger, 2020), and betas decompose into continuous, jump, and overnight components with distinct pricing (Bollerslev et al., 2016, 2020); jump-tail risk commands a large, time-varying premium (Bollerslev and Todorov, 2011). We deliberately use a single, strictly lagged market loading per stock-month rather than these richer structures, because any loading estimated on the event month would overfit the very minutes it is used to classify; the cost, a coarser residual, is absorbed by the placebo contrast.

2.11. Positioning and contribution

Each literature 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 and repeated on 2026-10-17—no prior study combines minute-level resolution, the full U.S. equity cross-section including delisted securities, and a 1992–2026 sample spanning floor trading, decimalization, Reg NMS fragmentation, and ETF-intermediated algorithmic markets. The joint analysis of market-wide and network-concentrated propagation, its secular evolution, and the evolution of synchronized extreme minutes at that resolution and length—with placebo-timing, placebo-network, announcement-exclusion, and tick-regime 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, which fixes the data used throughout. The archive carries no explicit licence; we use it for research with attribution and do not redistribute the minute data, only the aggregates, event-level estimates, and code. 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 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. Two properties of the source cannot be established from the data and are carried as limitations. The archive does not document whether bars aggregate consolidated-tape trades or a single venue; the presence of sparse low-volume prints after the 13:00 close on early-closing days suggests off-exchange trades are included, but volumes could not be reconciled against a consolidated benchmark for more than a few securities. And securities are identified by ticker alone, without a permanent identifier or security-type label, so that ETFs, ADRs, preferred shares, and warrants are present and can be screened only heuristically.

3.2. Full-sample audit

The 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, which rules out any daylight-saving mishandling (an offset error would displace roughly 15% of session rows in March and November). Early-closing sessions are truncated at the true exchange close. Integrity failures are negligible: 100 impossible OHLC bars (concentrated in 1993–1995, flagged and excluded) and 10,128 duplicate ticker–minute keys (1.3 per million), which we merge under a logged deterministic rule. No bar has a non-positive price or a zero or negative volume. A pilot audit of three months (June 1992, June 2009, December 2025) conducted before the full download, summarized in the Online Appendix, established the same facts on 58.8 million bars and verified volume units against public daily volumes for benchmark securities.

The panel spans 81,306 distinct tickers forming 221,957 continuous spells, a spell being a run of trading days without a gap longer than thirty calendar days; 29,112 tickers have multiple spells, so spells, not tickers, are the identity unit of every statistic that follows a security through time. Entries and exits occur continuously, with exits running between 800 and 1,600 per year: delisted securities are present throughout, and only 22–26% of the tickers active in June 1992 appear in 2009 or 2025. The data are therefore not a backfilled survivor universe, although we cannot verify that every delisting is captured to its final trading day. There is one discrete source-coverage expansion in 2020 (about 11,000 entries, largely over-the-counter names) that fixed-rule lagged universes absorb (Figure 1). A conservative overnight-ratio screen (canonical split ratios within $\pm 1.5\%$, both-day dollar volume of at least \$50,000, prior close of at least \$1) 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. Dividends and other distributions are not detected; their overnight effect never enters an intraday return because returns are computed within the trading day only. Table 1 tracks the sample from raw bars to the regular-hours analysis panel.

Table 1. 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 1. 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.

3.3. Universes

Coverage widens from roughly 8,300 symbols per month in 1992 to about 23,100 in late 2025, and much of the recent expansion reflects over-the-counter instruments rather than listed common equity. All analyses therefore operate on universes defined by strictly lagged rules: a security is eligible in month m if in month $m-1$ it traded on at least ten days with a median close of at least \$5 and a median of at least 100 regular-hours bars per day, and the top-500, top-1000, and top-3000 universes rank eligible securities by month- $(m-1)$ regular-hours dollar volume. Membership never uses information from month m . A ticker-shape heuristic (fifth-letter suffixes and warrant, unit, and preferred suffixes) flags likely non-common instruments; it does not identify ETFs, which trade under ordinary tickers.

3.4. Known limitations

The archive contains no quotes, spreads, depth, or trade direction; all statements about liquidity in this paper rest on dollar volume and the high–low range. Security types are unlabeled, prices are unadjusted, venue scope is undocumented, and the licence is unspecified. Pre-2001 minute returns are discrete at eighths and sixteenths, and roughly half of the one-minute returns of top-1000 securities in the mid-1990s are exactly zero (52% in 1996 against 13% in 2024, Online Appendix); any second-moment or tail statistic compared across the decimalization boundary inherits this discreteness. We flag the tick-regime caveat wherever it binds and anchor every secular comparison on a within-era placebo contrast.

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 construct 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 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 in nearly every minute from the mid-1990s onward (median 382–390 bars per day), 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; Lo and MacKinlay, 1990a) cannot masquerade as absent propagation. Even so, the number of receivers with an observable return at a given minute averages 237 in the 1990s and about 427 over the full sample, and every share below is computed over observed receivers only.

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 and $r_{m,t}$ without the superscript denotes the full equal-weighted mean. Without the leave-one-out correction a large own return mechanically moves the factor and attenuates the measured shock; with several hundred 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 $r_{m,t}$ over all minutes of month $m-1$ in which the market aggregate has at least 50 contributing names, provided security i has at least 300 valid return observations in that month; 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 stock-month is deliberately coarse: it cannot overfit the event minutes it will later be used to classify, and any misspecification it induces—including the downward bias in betas estimated from nonsynchronous minute returns (Scholes and Williams, 1977)—is common to real and placebo minutes of the same stock-month and is absorbed by the placebo contrast of 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. Let MAD_i be the median absolute deviation of $e_{i,\cdot}$ from its median over month $m-1$, and let $u_{i,t} = e_{i,t} / (1.4826 \text{MAD}_i)$. The minute-of-day factor is the cross-sectional median of $|u|$ at each minute of day τ in month $m-1$, normalized to unit mean over the session,

$$f_\tau = \frac{\text{med}_{(i,t): \tau(t)=\tau} |u_{i,t}|}{\frac{1}{390} \sum_{\tau'} \text{med}_{(i,t): \tau(t)=\tau'} |u_{i,t}|}, \quad z_{i,t} = \frac{e_{i,t}}{1.4826 \text{MAD}_i f_{\tau(t)}}. \quad (4)$$

The cross-sectional median makes the periodicity estimate robust to the jumps it is later used to detect, in the spirit of Boudt et al. (2011). The constant 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. A security whose prior-month residuals are exactly zero in more than half of the minutes has $\text{MAD}_i = 0$ and is excluded from the month, which acts as a staleness screen that binds mainly in the pre-decimalization years. The same

construction applied to $r_{m,t}$, 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 $|z_{i,t}| \geq 6$, with the sign recorded and analyzed separately. The threshold is deliberately extreme: it selects 1,854,384 events over 410 months, about 2.5 per stock-week in the top-500 universe, 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. The first requires a calm market factor, $|z_{m,t}| < 1$, so that the event does not coincide with a market-wide movement. The second requires that no simultaneous wave be under way: the share of observable securities with $|z_{j,t}| > 4$ at minute t is at most 2%, and at least 50 securities are observable. Both filters are measured, not assumed: minutes excluded by the second filter 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. The filters do not, however, require that the source be the only $|z| \geq 6$ event in its minute: with a 2% cap on simultaneous $|z| > 4$ names, up to about nine other securities may move sharply at t . In practice 22% of events are alone in their minute and the mean number of source events per event minute is 2.1; Section 6 reports every estimate separately for singleton minutes.

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. With

N_{t+h} the number of observable receivers, the outcomes are

$$\text{Breadth}_{i,t}(h) = \frac{1}{N_{t+h} - 1} \sum_{j \neq i} \mathbf{1}\{|z_{j,t+h}| > 2\}, \quad (5)$$

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

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

Breadth counts receivers with a *tail exceedance*, a standardized residual beyond two in absolute value; we avoid the word “jump” for this object, since the threshold is far below any formal jump-test cutoff. The threshold is calibrated to the fat-tailed base rate: in placebo windows about 0.10 of receivers exceed it, high enough that breadth is a well-populated share whose changes are measured with precision, low enough to respond to genuine tail activity. Intensity captures diffuse sub-threshold responses, and the signed outcome captures directional comovement with the source. 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 moves at $t + 3$ and is quiet at $t + 15$ still counts; its base rate is mechanically higher (about 58% at $h = 15$ in placebo windows) 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 is the differencing baseline: shocks arrive disproportionately in minutes of elevated ambient activity, and the post-minus-pre difference removes that level. It is also informative in its own right, and Section 6 reports the pre-event level of each arm rather than asserting their equality.

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 e over month $m - 1$ (pairwise complete observations; a neighbor set is used only if at least ten of its members are present in month m). Splitting equation (5) into

$$\text{Breadth}_{i,t}^{\mathcal{N}}(h) = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \mathbf{1}\{|z_{j,t+h}| > 2\}, \quad (8)$$

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

yields the *network contrast* $\text{Breadth}^{\mathcal{N}} - \text{Breadth}^{\overline{\mathcal{N}}}$, 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 about twenty. Because neighbor sets are lagged, the contrast cannot reflect correlations estimated on the event itself; and because the identical contrast is computed for placebo events, any unconditional tendency of correlated stocks to be more or less active than the rest of the universe nets out in the real-minus-placebo *network differential* of Section 5.

Two properties of this construction must be kept in view. First, neighbors are identified from one-minute residual correlations, and the precision of that identification is not constant over time: when the Epps effect is strong, as it was in the 1990s (Section 6), one-minute correlations are dominated by asynchronous trading and discreteness, and the twenty “closest” names are a noisier selection of true economic peers than they are in the 2010s. Secular changes in the network differential therefore combine changes in transmission with changes in how well the network is measured. Second, the network contrast is computed at $t + h$ only; its pre-event counterpart at $t - 5$ was not stored by the pipeline and is listed among the analyses left for future work.

4.6. Systemic co-jump minutes

A minute is a *systemic co-jump minute* if the share of observable securities with $|z_{j,t}| > 4$ exceeds 2% (with at least 50 observable names)—exactly the complement of the second isolation filter. 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 observable 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. The same discreteness also inflates the severe-minute count before 2001, and we do not compare the 1990s with later decades on this series.

4.7. What these measures can and cannot see

Three mechanical properties discipline interpretation. First, breadth responds to tail activity only; diffuse sub-threshold responses appear in intensity, which is reported alongside breadth. 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 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 is reported for every headline estimate.

5. Empirical design

5.1. The estimand

For an outcome Y from Section 4.4, define the within-event change $\Delta Y_{i,t}(h) = Y_{i,t}(h) - Y_{i,t}(-5)$, where $Y_{i,t}(-5)$ denotes the same outcome evaluated five minutes before the event. The baseline estimand is the difference-in-differences

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

estimated as the difference between the two arms' sample means. In potential-outcomes language, the placebo arm estimates the path of cross-sectional activity that the event window would have followed had the source stock not moved, 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.

The inner difference has a cost that we quantify rather than assume away. If real events select minutes whose ambient activity is elevated for transient reasons, part of that elevation reverts over the following minutes, and $\Delta\Delta$ is biased downward by the reversion. We therefore report, alongside (10), the post-event level difference without pre-differencing and a stratified estimator that compares real and placebo post-event levels within deciles of the pre-event level,

$$\hat{\theta}^S(h) = \sum_{b=1}^{10} w_b \left(\bar{Y}_b^{\text{real}}(h) - \bar{Y}_b^{\text{placebo}}(h) \right), \quad w_b = \frac{n_b^{\text{real}}}{\sum_{b'} n_{b'}^{\text{real}}}, \quad (11)$$

where b indexes deciles of $Y_{i,t}(-5)$ over the pooled sample. Because the placebo arm has a lower latent activity level than the real arm within any stratum defined on a noisy pre-event measurement, regression to the mean pushes (11) upward; the two estimators therefore bracket the market-wide response from below and from above.

5.2. Placebo construction

For each real event (i, t) the matched placebo is drawn as follows, with a fixed seed recorded in the replication package. The candidate set collects all minutes of the same month with the same minute of day $\tau(t)$ on a different trading day at which security i has an observable return with $|z_{i,\cdot}| < 2$, the market-calm filter holds, and the simultaneity filter holds. One candidate is drawn uniformly at random, and every outcome of Section 4.4 is computed at the placebo minute by the identical code path. The placebo is labelled with the real event’s magnitude and sign so that all conditional cuts—by $|z|$ bin, by sign, by number of co-arriving events—compare real and placebo events within the same cell. Events with an empty candidate set contribute to the real arm only; the match rate is 95.2% over the full sample, 71% in the 1990s, when calm candidate minutes with an observable return are scarcer, and at least 96% in each later decade. The placebo is not matched on the pre-event level of the outcome, which is why the pre-event levels of the two arms can differ and are reported.

5.3. Inference

Outcomes are correlated across events within a trading day—mechanically so, because events of the same minute share receivers and 2.1 source events on average share each event minute—and within a source security across days. All inference therefore resamples clusters: 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. Standard errors are bootstrap standard deviations. Two alternatives are reported: a month-cluster variant (resampling year-months), which roughly doubles the standard error of the network differential without changing any conclusion, and a ticker-cluster variant (resampling source securities), which more than doubles it. With 7,282 event days and 1.85 million events, the minimum detectable effect for market-wide breadth at $h = 5$ at the 5% level with 80% power is about 2.5×10^{-4} , or one additional responding security per 9 events given 427 observable receivers on average. Null results below are powered nulls in this sense.

5.4. Identification: what is and is not claimed

The design identifies the predictive association between an isolated idiosyncratic extreme and subsequent cross-sectional tail activity, purged of intraday seasonality, calm-market selection, and the stock’s own base rates, and bounded with respect to ambient-activity selection as described above. 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 repricing linked firms on the same news) 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 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 and diagnostic battery

Five checks are reported. (1) *Pre-event levels*: the levels of every outcome at $t - 5$ in the two arms, with their difference; a non-zero difference is not a failure of the design but a measured quantity that the estimators above handle in different ways. (2) *Placebo network contrast*: the neighbor-minus-other contrast computed on placebo events alone, over the full sample and by decade. (3) *Announcement windows*: excluding the 10:00 and 14:00–14:30 ET windows that host most scheduled releases. (4) *Coverage*: fixed-rule lagged universes make the 2020 source-coverage expansion irrelevant to any comparison. (5) *Tick regime*: the post-decimalization subsample is reported for every headline estimate. Tables 3, 5, and 6 quantify each.

6. Results

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

Figures 2 through 4 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 2a), 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 3). A large part of the 1990s level is bid–ask bounce and price discreteness under eighth and sixteenth ticks rather than volatility (Roll, 1984; Harris, 1991): a single tick on a \$20 stock is a 0.6% return, and about half of all one-minute returns in the mid-1990s are exactly zero (Online Appendix). The decline is nonetheless not only a tick artifact. Measured from 2002, the first full year after decimalization, the frequency of $|r| > 1\%$ minutes falls from 50 per 10,000 to 4.8 in the 2010s, roughly tenfold, and the cross-section of daily realized variance shifts down over the same years (Online Appendix).

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 at a smaller scale (Figure 3). Part of the 2020–2021 level reflects the composition of the top-1000 universe in those years, which dollar-volume ranking populates with leveraged exchange-traded




Figure 2. 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.

products and distressed issuers that we cannot label; the qualitative pattern—extremes became rarer but, when the market is stressed, arrive in synchronized bursts—survives the exclusion of heuristically identified non-common tickers and is the subject of Section 7.

Third, cross-sectional comovement at the one-minute horizon rose from nothing to most of its half-hour level. 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 was invisible at the minute, 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 at a lower level (0.21 against 0.24; Figure 4). Two mechanisms produce this flattening and our previous-tick synchronized grids do not separate them. One is mechanical: nonsynchronous trading and stale prices attenuate measured high-frequency correlation (Renò, 2003; Hayashi and Yoshida, 2005; Zhang, 2011), and the share of stale one-minute prices fell from 52% to 13% between 1996 and 2024. The other is economic: common information is now impounded across the cross-section within the minute. Tóth and Kertész (2009) reach the same conclusion for 1993–2003 and find that denser trading explains only part of the change. Whatever the mix, the operational consequence for this paper is the same: one-minute residual correlations identify economic neighbors far more sharply after 2005 than before, a point we carry into the interpretation of the network results.

6.2. Baseline: isolated shocks and market-wide tail activity

Table 2 reports the baseline estimates of equation (10) for 1,854,384 isolated $|z| \geq 6$ source events and 1,766,058 matched placebos. The difference-in-differences breadth response is economically negligible at every horizon: at $h = 5$ minutes the point estimate is -1.04×10^{-4} (bootstrap standard




Figure 3. 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 largely reflect eighth/sixteenth tick bid–ask bounce.

error 0.90×10^{-4}), and at $h = 30$ minutes -3.10×10^{-4} (1.01×10^{-4}). Against a placebo base rate of about 0.10 (Table 3, Panel A), the largest estimate in absolute value implies that an isolated six-sigma move changes the probability that any given other security shows a tail exceedance by about a third of one percent of the base rate, with a negative sign that is statistically distinguishable from zero at $h = 30$. Difference-in-differences intensity estimates are of the same order and alternate in sign.

Table 2. 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.

The design diagnostics of Table 3 qualify this null in three ways. Panel A shows that real events do not arrive in ordinary minutes: five minutes before the event, breadth is $1,044 \times 10^{-4}$ for real events against $1,004 \times 10^{-4}$ for placebos, a difference of 40.3×10^{-4} (standard error 1.2), and intensity is higher by 0.074 (0.012). A six-sigma idiosyncratic move is thus more likely on days and in minutes when the whole cross-section is already active, even though the market factor is calm by construction. This is exactly the ambient level that the inner difference of (10) removes, and it is the




Figure 4. The flattening of the Epps curve. 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. Correlations are computed on calendar-clock grids and are therefore attenuated by nonsynchronous trading at short intervals, most strongly in 1996.

reason the difference-in-differences, not a post-event comparison, is the baseline; but it also means that the “no pre-event response” statement sometimes made for designs of this kind does not hold here as a statement about levels, and we do not make it.

Panel C reports the three estimators side by side. Without pre-differencing, real events show 40×10^{-4} more breadth than placebos at $h = 5$, almost exactly the pre-event gap, so that the difference-in-differences is close to zero. Within deciles of the pre-event level, the gap is 19×10^{-4} (0.9) at $h = 5$ and 18×10^{-4} (1.0) at $h = 30$. As argued in Section 5, the two estimators bracket the market-wide response: between zero and roughly 2% of the base rate. Either bound is small relative to the network response of Section 7, and the conclusion we draw is correspondingly bounded: isolated idiosyncratic extremes raise market-wide tail activity by at most a few percent of its base rate within two hours.

Panel B separates events by the number of source events detected in the same minute. For the 22% of events that are alone in their minute, the difference-in-differences is 1.7×10^{-4} (1.7) at $h = 5$ and 3.7×10^{-4} (1.9) at $h = 30$: slightly positive and within two standard errors of zero. The negative full-sample estimates at longer horizons come from minutes with three to five co-arriving events (-8.3×10^{-4} , 2.1, at $h = 30$), which are the minutes with the most elevated ambient activity and the strongest reversion. The negative sign of the pooled estimate is therefore a mean-reversion artifact of co-arrival minutes, not evidence that shocks calm the market.

Three further features of Table 2 deserve emphasis. First, this is a powered null: the minimum detectable difference-in-differences at $h = 5$ is about 2.5×10^{-4} , one additional responding security per 9 events. Second, the null is uniform in time: Figure 5 shows the same flat response profile




Figure 5. Difference-in-differences breadth by decade. Point estimates of $\Delta\Delta$ breadth (equation (10)) with 95% day-block bootstrap bands, by horizon and decade.

in each of the four decades, and the Online Appendix shows it year by year, 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. Third, Panel D of Table 3 shows that the standard errors are insensitive to the clustering unit for the market-wide statistics.

Conditioning does not change the picture (Figure 6): estimates by shock sign and by calm versus stressed market state are within two standard errors of zero. Across magnitude bins the point estimates rise from significantly negative for 6σ – 8σ events, which are the most numerous and the most affected by the co-arrival reversion just described, to a positive but insignificant 3.47×10^{-4} (2.95) for 12σ – 20σ events. The signed response of Panel E is noisy and never significant: the cross-section does not drift in the direction of the source shock. None of this resurrects market-wide propagation; it motivates asking where a concentrated response would hide, which is the subject of the next section.

7. Mechanisms

The baseline bounds the market-wide response to isolated idiosyncratic extremes at a few percent of the base rate. This section asks where transmission lives if not there: in the shocked stock’s network neighborhood, 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

Market-wide breadth spreads any response across several hundred receivers; a transmission channel operating through economic or statistical linkages concentrates it among a handful. For every event

Table 3. Design diagnostics: pre-event levels, co-arrivals, mean reversion, and clustering

<i>Panel A: Pre-event levels at $t - 5$ ($\times 10^{-4}$)</i>			
	Real	Placebo	Real – placebo
Breadth	1,044	1,004	40.25 (1.18)
Intensity	13,226	12,488	737.77 (115.12)
<i>Panel B: By number of source events in the same minute</i>			
Events in minute (share of events)	$\Delta\Delta\text{Breadth}_5$	$\Delta\Delta\text{Breadth}_{30}$	Network diff ₅
1 (21.7%)	1.74 (1.67)	3.69 (1.89)	51.01 (2.55)
2 (27.1%)	-2.97 (1.80)	-2.29 (1.98)	64.17 (2.67)
3–5 (43.6%)	-3.01 (1.79)	-8.35 (2.13)	66.87 (2.99)
6+ (7.7%)	9.12 (5.19)	1.39 (5.68)	76.49 (7.40)
<i>Panel C: Market-wide breadth, alternative estimands ($\times 10^{-4}$)</i>			
Horizon	Post level, real – placebo	$\Delta\Delta$ (baseline)	Pre-level stratified
$h = 5$	39.63 (1.28)	-1.04 (0.87)	19.48 (0.88)
$h = 15$	39.46 (1.36)	-1.46 (1.00)	19.59 (0.98)
$h = 30$	37.93 (1.42)	-3.10 (1.07)	18.41 (1.04)
<i>Panel D: Bootstrap standard errors by clustering unit ($\times 10^{-4}$)</i>			
Statistic	Estimate	Day clusters	Ticker clusters
$\Delta\Delta\text{Breadth}_5$	-1.04	(0.90)	(0.96)
$\Delta\Delta\text{Breadth}_{30}$	-3.10	(1.06)	(0.93)
Network diff ₅	64.02	(1.89)	(4.72)
<i>Panel E: Signed response, real – placebo ($\times 10^{-4}$)</i>			
Horizon	$h = 1$	$h = 5$	$h = 30$
$\text{sign}(e_{i,t}) \times \text{mean } z_{j,t+h}$	151.33 (104.90)	9.57 (91.95)	-65.70 (89.23)

Notes: All statistics from the baseline event files (isolated $|z| \geq 6$ source shocks, lagged top-500 universe, 1992–2026; matched within-stock placebos). Day-block bootstrap standard errors (500 replications) in parentheses unless stated. Panel A: levels of the outcomes five minutes before the event minute. Panel B: events classified by the number of source events detected in the same minute (the isolation filters cap the share of simultaneous $|z| > 4$ names at 2%, not the number of $|z| \geq 6$ events); each placebo inherits the class of its real event. Panel C: ‘post level’ is the real-minus-placebo difference in breadth at $t + h$ without pre-differencing; ‘pre-level stratified’ reweights the same difference within deciles of the pre-event breadth level. Panel D: standard errors from resampling trading days versus resampling source tickers. Panel E: cross-sectional mean standardized return of receivers signed by the source shock.




Figure 6. 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.

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 network differential is large relative to the market-wide estimates, precisely estimated, and decays slowly with the horizon: five minutes after an isolated $|z| \geq 6$ move, the source's neighbors are 64.0×10^{-4} more likely than other receivers to show a tail exceedance, net of the same contrast for placebos, with a day-block bootstrap standard error of 1.8×10^{-4} and a ticker-cluster standard error of 4.7×10^{-4} ; the differential declines to 55.4×10^{-4} at fifteen minutes and 49.6×10^{-4} at thirty. Against a neighbor base rate of about 0.10, a neighbor's exceedance probability rises by roughly six percent of its unconditional level. The response is symmetric in the sign of the shock (66×10^{-4} after negative and 63×10^{-4} after positive moves). Idiosyncratic extremes do propagate—to the connected few, not to the diffuse many, which is why market-wide breadth cannot detect them.

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.

Table 5 decomposes the differential into its four components. Over the full sample the contrast is driven by the real arm: neighbors of a shocked stock show about 61×10^{-4} more exceedances than other receivers, while the placebo contrast is -2.7×10^{-4} . By decade, however, both arms move. In the 1990s, neighbors show 228×10^{-4} more exceedances than other receivers after a real event and 200×10^{-4} more at placebo minutes; in the 2010s the real contrast is 37×10^{-4} and the placebo contrast is -44×10^{-4} . The placebo contrast is thus far from zero within decades even though it averages to nearly zero over the sample: in the 1990s, the twenty most correlated names are simply the most actively traded names, which show more exceedances in every minute; after 2000, the neighbors of a stock that is quiet at the placebo minute are themselves quieter than the rest of the universe, as correlation implies. The network differential is therefore legitimately a difference of two contrasts, and only the real-minus-placebo column of Table 5 carries a propagation interpretation. Panel B of Table 3 adds that the differential survives the restriction to minutes in which the source is the only $|z| \geq 6$ event (51×10^{-4} , 2.5), so it is not an artifact of correlated pairs moving in the same minute and remaining active five minutes later, although co-arrivals do account for about a fifth of the pooled estimate.

Table 5. Components of the network differential at $h = 5$, by decade

Sample	Real events			Placebo events			Real – placebo
	Neighbors	Others	Diff.	Neighbors	Others	Diff.	
Full sample	1,103	1,043	61.3 (3.4)	1,004	1,006	-2.7 (3.1)	64.0 (1.8)
1992–1999	1,191	1,006	228.3 (7.7)	1,121	932	200.0 (7.9)	28.3 (6.9)
2000–2009	1,099	1,110	-10.7 (4.3)	1,036	1,076	-40.6 (4.1)	30.0 (2.2)
2010–2019	1,060	1,023	37.0 (6.2)	941	985	-43.7 (5.2)	80.6 (3.7)
2020–2026	1,144	987	157.1 (7.2)	1,023	955	67.5 (5.7)	89.6 (4.3)

Notes: Share of receivers with $|z| > 2$ five minutes after the event minute, among the source’s twenty lagged-correlation neighbors and among all other receivers, for real events and for their matched placebos ($\times 10^{-4}$). ‘Diff.’ is neighbors minus others within an arm; the last column is the network differential of Table 4 (real minus placebo). Day-block bootstrap SEs (500 replications) in parentheses.

The evolution across decades (Figure 7) is the paper’s sharpest dynamic fact and its most delicate one. The five-minute network differential was 28×10^{-4} in 1992–1999 and 30×10^{-4} in 2000–2009, then rises to 81×10^{-4} in 2010–2019 and 90×10^{-4} in 2020–2026, a threefold increase that coincides with the era of algorithmic liquidity provision, correlation-aware execution, and index intermediation, and that occurs while market-wide spillover stays bounded near zero. Two readings are consistent with this pattern and the data do not separate them. Under the first, transmission along economic linkages became stronger and faster—basket arbitrage and statistical-arbitrage strategies convert a move in one name into trades in its peers within minutes, as Ben-David et al. (2018) and Foucault et al. (2016) would predict. Under the second, transmission was always




Figure 7. Network-concentrated propagation by decade. Neighbor minus other-receiver breadth at $h \in \{5, 15, 30\}$, real minus placebo, with 95% day-block bootstrap intervals. Neighbors are the source's top-20 absolute-correlation partners estimated on prior-month one-minute residuals.

present but the one-minute correlations from which we identify neighbors were too noisy in the 1990s to find the true peers: with the Epps effect of Figure 4, the prior-month one-minute residual correlation of a genuine industry pair in 1996 is barely distinguishable from sampling noise, and a neighbor set contaminated by false peers dilutes the contrast mechanically. The decline of the 1990s differential at $h = 5$ relative to $h = 30$ (Table 4), the opposite of the decay seen in later decades, is consistent with the second reading, since slow comovement is what an Epps-era network would transmit. The honest summary is that propagation along the measured network at least tripled, and that transmission became sharper rather than broader; how much of the increase is a change in the market and how much a change in the measurement requires neighbor sets built on coarser-sampled or transaction-time correlation estimators (Hayashi and Yoshida, 2005), an extension we leave for future work.

7.2. Cumulative responses

Point-in-time breadth at $t + h$ misses a receiver that moves 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} , standard error 1.9×10^{-4} , against a base rate near 58% at fifteen minutes). The market-wide bound is therefore not an artifact of sampling the response at isolated instants; transient market-wide cascades after isolated idiosyncratic extremes are absent as well.

7.3. Systemic co-jump minutes

The isolation filters exclude minutes in which more than 2% of the universe moves beyond four robust standard deviations simultaneously. These minutes are the natural home of the fragility hypothesis, and their time series (Figure 8) delivers two messages, each with a comparability limit. First, the ordinary cross-section de-synchronized secularly: with idiosyncratic tail activity two orders of magnitude rarer, minutes breaching the fixed 2% threshold fell from near-ubiquity in the sparse, wide-tick 1990s to about 131 per 1,000, roughly one minute in eight, in 2021–2025. Under the same standardization the early 1990s also produce 128 severe minutes per 1,000 and a mean affected share of 7.1% within co-jump minutes—both above 2008—which we read as discreteness rather than fragility: when most minute returns are zero or one tick, a two-tick move in many names at once breaches a four-MAD threshold without any economic synchronization. The fixed-threshold count and the pre-2001 severity series are therefore not fragility measures, and no claim in this paper compares synchronization across the decimalization boundary.

Second, within the post-decimalization era, conditional severity moves sharply 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 moving together) rise from 9.53 per 1,000 in 2017–2019 to 83.7 per 1,000 in 2020, with 91.1 per 1,000 in 2008 against fewer than 6 per 1,000 in 2006 and 2019: an increase of roughly an order of magnitude in both crises. Combined with the evidence of Figure 3 that $|r| > 5\%$ minutes rose by a factor of 391 in 2020, the pattern is consistent across measures and with the collective-synchronization events documented by [Calcagnile et al. \(2018\)](#) on a shorter sample: modern stress arrives as market-wide synchronized minutes, not as single-name events. Whether those minutes are common news arrivals or endogenous synchronization of the kind modelled by [Bormetti et al. \(2015\)](#) is not identified by bars alone.

Taken together, the mechanism evidence sharpens the baseline into a bounded statement: idiosyncratic extremes propagate narrowly along the lagged correlation network, that measured channel at least tripled in the algorithmic era, and market-wide tail risk in the post-decimalization era is the province of common-shock minutes whose severity in crises, not their frequency in calm times, is the binding object.

8. Robustness

Table 6 re-estimates the market-wide difference-in-differences at $h \in \{5, 30\}$ and the network differential at $h = 5$ under eleven designs; Tables 3 and 5 report the design diagnostics discussed above. We describe only checks that were executed; analyses that the replication package does not contain are listed as such in the Online Appendix and in the concluding section.




Figure 8. Systemic co-jump minutes, 1992–2026. (a) Minutes per 1,000 in which more than 10% of observable top-500 names exceed $|z| > 4$ simultaneously (log scale). (b) Mean affected share within co-jump minutes ($> 2\%$ threshold). Pre-2001 levels are inflated by price discreteness and are not comparable with later years.

Tick regime. Pre-2001 minute returns embed eighth- and sixteenth-tick bid–ask bounce. Restricting events to 2002 onward leaves the market-wide estimates unchanged (-1.03 and -3.04×10^{-4}) and the network differential at 65.8×10^{-4} . *Time of day.* Restricting events to 10:00–15:30 ET, which removes the opening and closing periods where seasonal scaling is most demanding, and separately excluding the 10:00 and 14:00–14:30 ET windows that host most scheduled macroeconomic releases, changes no estimate by more than one standard error. *Market state.* Excluding 2008–2009 and 2020 leaves both objects unchanged; within the crisis years alone the market-wide estimate is imprecise ($2.7 \pm 5.1 \times 10^{-4}$) and the network differential is, if anything, larger (68×10^{-4}). *Source magnitude.* Raising the source threshold to $|z| \geq 8$ and $|z| \geq 12$ moves the market-wide point estimates toward zero or slightly positive, consistent with the co-arrival pattern of Table 3, and lowers the network differential from 64 to 50 and 34×10^{-4} . The decline with magnitude is the one systematic pattern in the table; it is consistent with the largest moves being firm-specific news that peers correctly do not price, and also with the largest moves being the ones most often accompanied by halts or erroneous prints that bars cannot flag. *Universe composition.* Excluding source securities flagged as likely non-common by the ticker-shape heuristic leaves every estimate within rounding of the baseline; ETFs cannot be identified and are not excluded. *Inference.* Month-cluster resampling raises the standard error of the network differential from 1.8 to 3.4×10^{-4} and ticker-cluster resampling to 4.7×10^{-4} , leaving a t -statistic above fourteen in every case; the market-wide standard errors are insensitive to the clustering unit. *Placebo network.* The neighbor-minus-other contrast computed on placebo events alone is $-2.7 \pm 2.8 \times 10^{-4}$ over the full sample. As Table 5 shows, this average conceals decade-level offsets of opposite signs, which is why every network statistic in the paper is a real-minus-placebo difference and never a real-arm contrast on its own.

Table 6. 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).

Several checks that a complete treatment would include were not executed and are not claimed. The pipeline stores events at the $|z| \geq 6$ threshold and the receiver threshold $|z| > 2$ only, so lower source thresholds ($|z| \geq 4, 5$), a higher receiver threshold ($|z| > 3$), and the top-1000 and top-3000 universes would require re-running the detection stage on the raw panel. Neighbor sets built on 5- or 30-minute residual correlations, or on a transaction-time covariance estimator, and the pre-event ($t - 5$) network contrast, are the two extensions most needed to settle the interpretation of the secular increase in the network differential. Permuted-network placebos that preserve industry and liquidity composition, a separate estimate of propagation inside announcement windows, and a Hawkes-type model of the co-jump minute series are natural further steps.

9. Conclusion

Three decades of minute-level data on the full U.S. equity cross-section allow the title’s question to be answered in its three parts. Faster: common information that in the mid-1990s took tens of minutes to appear in cross-stock comovement now appears within the minute, whether because incorporation accelerated or because the mechanical attenuation of nonsynchronous trading receded, and most likely both. Wider: no. Isolated idiosyncratic extremes—six-sigma one-minute moves in the largest five hundred securities—raise market-wide tail activity by at most a few percent of its base rate within two hours, in every decade and for either sign; what transmission exists is concentrated among the shocked stock’s twenty closest lagged-correlation neighbors, and that

concentrated channel became roughly three times larger between the floor-trading and algorithmic eras, so transmission became sharper rather than broader. More fragile: only in a bounded sense. Within the post-decimalization era, ordinary extreme minutes became an order of magnitude rarer while minutes in which a tenth of the universe moves together are roughly ten times more frequent in 2008 and 2020 than in calm years; across the decimalization boundary the question cannot be answered with threshold-based measures, because price discreteness manufactures synchronization in the 1990s.

The evidence is deliberately associational where identification runs out, and the paper reports the diagnostics that qualify its own headline numbers. Real events arrive in minutes of elevated ambient activity, so the market-wide response is bracketed rather than pinned at zero; the network differential is a difference between two non-zero contrasts, and its secular increase combines a change in transmission with a change in how precisely one-minute correlations identify economic neighbors; and the co-jump series conflates common news with endogenous synchronization. OHLCV bars do not observe quotes, depth, order flow, or security type, so liquidity operates through proxies and ETFs cannot be isolated; tickers are not permanent identities, so long-run statements run within screened spells. The extensions that would most sharpen the conclusions are, in order, neighbor sets estimated on coarser-sampled or transaction-time correlations, the pre-event network contrast, lower source and higher receiver thresholds on wider universes, and the attribution of co-jump minutes to news using timestamped announcement data.

For market design and risk monitoring the practical implication is direct. Surveillance aimed at single-name events monitors the part of the system that modern markets absorb, or transmit only to a handful of linked names; the binding constraint for market-wide stability is the market-wide minute, and the margin that moved over three decades is the sharpness of transmission along statistical linkages, not its breadth.

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.
- , **Allen Carrion, Laura Tuttle, and Kumar Venkataraman**, “Liquidity, resiliency and market quality around predictable trades: Theory and evidence,” *Journal of Financial Economics*, 2016, 121 (1), 142–166.
- 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 Loriana 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, and Yuan Xue**, “Volume, Volatility, and Public News Announcements,” *Review of Economic Studies*, 2018, 85 (4), 2005–2041.
- , —, **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.
- , **Viktor Todorov, and Sophia Zhengzi Li**, “Jump tails, extreme dependencies, and the distribution of stock returns,” *Journal of Econometrics*, 2013, 172 (2), 307–324.
- Bormetti, Giacomo, Lucio Maria Calcagnile, Michele Treccani, Fulvio Corsi, Stefano Marmi, and Fabrizio Lillo**, “Modelling systemic price cojumps with Hawkes factor models,” *Quantitative Finance*, 2015, 15 (7), 1137–1156.
- 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.
- , **Christophe Croux, and Sébastien Laurent**, “Robust estimation of intraweek periodicity in volatility and jump detection,” *Journal of Empirical Finance*, 2011, 18 (2), 353–367.
- 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.
- Calcagnile, Lucio Maria, Giacomo Bormetti, Michele Treccani, Stefano Marmi, and Fabrizio Lillo**, “Collective synchronization and high frequency systemic instabilities in financial markets,” *Quantitative Finance*, 2018, 18 (2), 237–247.
- 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.
- , **Roel Oomen, and Roberto Renò**, “The drift burst hypothesis,” *Journal of Econometrics*, 2022, 227 (2), 461–497.
- 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.
- Dungey, Mardi, Deniz Erdemlioglu, Marius Matei, and Xiye Yang**, “Testing for mutually exciting jumps and financial flights in high frequency data,” *Journal of Econometrics*, 2018, 202 (1), 18–44.
- Easley, David, 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.
- Gilder, Dudley, Mark B. Shackleton, and Stephen J. Taylor**, “Cojumps in stock prices: Empirical evidence,” *Journal of Banking & Finance*, 2014, 40, 443–459.
- 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.
- 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.
- Hansen, Peter R. and Asger Lunde**, “Realized Variance and Market Microstructure Noise,” *Journal of Business & Economic Statistics*, 2006, 24 (2), 127–161.

- Harris, Lawrence**, “Stock Price Clustering and Discreteness,” *Review of Financial Studies*, 1991, 4 (3), 389–415.
- Hasbrouck, Joel 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.
- Hayashi, Takaki and Nakahiro Yoshida**, “On covariance estimation of non-synchronously observed diffusion processes,” *Bernoulli*, 2005, 11 (2), 359–379.
- 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.
- Jacod, Jean and Viktor Todorov**, “Testing for common arrivals of jumps for discretely observed multidimensional processes,” *Annals of Statistics*, 2009, 37 (4), 1792–1838.
- Johnson, Neil, Guannan Zhao, Eric Hunsader, Hong Qi, Nicholas Johnson, Jing Meng, and Brian Tivnan**, “Abrupt rise of new machine ecology beyond human response time,” *Scientific Reports*, 2013, 3, 2627.

- 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**, “An econometric analysis of nonsynchronous trading,” *Journal of Econometrics*, 1990, 45 (1–2), 181–211.
- and —, “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.
- Renò, Roberto**, “A closer look at the Epps effect,” *International Journal of Theoretical and Applied Finance*, 2003, 6 (1), 87–102.
- 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.
- Scholes, Myron and Joseph Williams**, “Estimating betas from nonsynchronous data,” *Journal of Financial Economics*, 1977, 5 (3), 309–327.
- Tóth, Bence and János Kertész**, “The Epps effect revisited,” *Quantitative Finance*, 2009, 9 (7), 793–802.
- 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.

Zhang, Lan, “Estimating covariation: Epps effect, microstructure noise,” *Journal of Econometrics*, 2011, 160 (1), 33–47.