

Pre-Trade Uncertainty and Liquidity Stress: The Timing of Quote Revisions

Ayush Jha*

Texas Tech University

Ali Jaffri†

North Dakota State University

Svetlozar Rachev‡

Texas Tech University

Frank Fabozzi§

Johns Hopkins University

September 21, 2026

Abstract

We study whether the timing of quote revisions contains information about liquidity stress that is not captured by the magnitude of price movements. We develop a subordinated-Lévy framework that separates return variation into a diffusion component and a component associated with irregular business time. The clock-dispersion parameter measures how unevenly price revisions occur over calendar time. Using S&P 500 index quotes around stress and policy events and 779 trading days of SPY National Best Bid and Offer data from 2007–2025, we find that clock dispersion is related to the dispersion of mid-quote revisions and rises sharply around FOMC statements. The response is concentrated in the statement window, while changes are not observed around CPI and employment releases. We find little evidence that clock dispersion forecasts liquidity conditions. The results support the interpretation that timing irregularity can help describe liquidity stress within an event window, while its forecasting role remains limited.

Keywords: Subordinated Lévy Processes; Stochastic Business Time; Normal Inverse Gaussian; Identification; Bid–Ask Quotes; Market Liquidity; Financial Stability.

JEL Classification: C32; G01; G12; G14; G28.

*Corresponding author. Department of Mathematics and Statistics, Texas Tech University, Lubbock, TX 79409. Email: ayush.jha@ttu.edu.

†College of Business, North Dakota State University, Fargo, ND 58105. Email: ali.jaffri@ndsu.edu.

‡Department of Mathematics and Statistics, Texas Tech University, Lubbock, TX 79409. Email: zari.rachev@ttu.edu.

§Carey Business School, Johns Hopkins University, Baltimore, MD 21202. Email: ffabozzi1@jhu.edu.

1 Introduction

Financial market stress is often summarized by the size of price movements, changes in quoted spreads, or measures of realized volatility. These measures are useful, but they do not describe when quote revisions occur. During periods of rapid information arrival, prices may be revised in short, concentrated bursts even when the aggregate increase in daily volatility is comparatively modest. The timing of these revisions is potentially relevant for understanding short-lived episodes of market liquidity stress, particularly around scheduled announcements.

This paper studies whether that timing dimension provides information about liquidity stress that is not contained in the magnitude of price movements. We develop a model-based measure of the irregularity of price revisions over calendar time. The framework treats market activity as stochastic business time and separates return variation into a diffusion component and a component associated with dispersion in that clock. The resulting clock-dispersion parameter, v , is our central empirical object. We do not interpret v as a measure of information arrival, liquidity risk, or a forecasting signal. Rather, we ask whether it provides a distinct description of how market prices adjust during episodes in which liquidity conditions change.

The empirical contribution is deliberately narrow. We document that clock dispersion is associated with the dispersion of mid-quote revisions and that it rises sharply during the statement window of FOMC announcements. The increase is much more concentrated in time than the corresponding change in daily volatility. This pattern is also visible relative to other announcement windows, although the comparisons do not identify causal effects because scheduled events occur at known times of day. The result suggests that the timing of quote revisions can provide information about the within-event evolution of market stress that is not conveyed by a single volatility measure.

The model supplies the measurement framework for this exercise. A doubly subordinated clock is not identified from return cumulants through fourth order, so we work with the identified single-subordinator case. This reduction yields a normal inverse Gaussian return distribution. Under the moment estimator, the model's uncertainty index is numerically equal to the sample standard deviation; consequently, the empirical content of the framework is not a new volatility measure. It is the decomposition of return variation and, in particular, the clock-dispersion component. This distinction is important because it prevents the index itself from being presented as a separate measure of volatility.

We use two empirical settings. The first examines one-minute S&P 500 index bid and ask quotes over 45 trading days surrounding major stress and policy episodes. These observations provide detailed within-day evidence but are limited in number. The second uses SPY National Best Bid and Offer data for 779 trading days from 2007 through 2025. The longer panel allows us to examine whether the fitted clock is related to observable quote revisions after accounting for day, year, and time-of-day effects and for conventional market conditions.

Several supporting results help delimit the central finding. Clock dispersion is positively related to the dispersion of mid-quote revision counts, with rank correlations around 0.77 in the panel. Around 166 FOMC statement days, clock dispersion in the statement window is substantially above its usual level, while realized volatility and quoted spreads provide a less concentrated description of the episode. We do not find a comparable response around

CPI or Employment Situation releases. At longer horizons, we find little evidence that clock dispersion predicts subsequent liquidity conditions. Finally, the aggregation behavior of the data is not fully consistent with a Lévy subordinator with independent increments. These results are important qualifications: they support the use of the clock as a descriptive measure of within-event timing, not as a general forecasting or early-warning indicator.

The paper adds to the market-liquidity and microstructure literature by examining a dimension of quote dynamics that is usually absorbed into aggregate measures of market activity. The time-change framework makes this dimension measurable in high-frequency data. The identification and aggregation results support the measurement framework rather than constituting separate empirical claims.

The remainder of the paper is organized as follows. Section 2 discusses the literature and the motivation for studying liquidity stress. Section 3 develops the model and measurement framework. Section 4 describes the data and empirical design. Section 5 presents the event evidence, and Section 6 examines the multi-year SPY panel. Section 7 discusses interpretation and limitations. Section 8 concludes.

2 Related Literature and Financial-Stability Motivation

Our paper connects three literatures: market microstructure and limit-order-book dynamics, time-changed and Lévy-process models in asset pricing, and uncertainty measurement.

The microstructure foundations rest on the information-based models of [Glosten and Milgrom \(1985\)](#) and [Kyle \(1985\)](#), which linked bid–ask spreads to inventory risk, adverse selection, and market depth. In our framework the stochastic clock measures how irregularly business time accrues, that is, how unevenly valuations are revised over calendar time. Empirical decompositions of the spread into adverse-selection, inventory-holding, and order-processing components ([Glosten and Harris, 1988](#); [Huang and Stoll, 1997](#)), together with information-share measures of price discovery ([Hasbrouck, 1991](#)), have been extended to high-frequency settings by [Biais et al. \(2005\)](#), [Hasbrouck \(2010\)](#), and [Aït-Sahalia and Jacod \(2010\)](#); comprehensive syntheses appear in [O’Hara \(1995\)](#) and [Madhavan \(2000\)](#). A parallel literature documents the statistical structure of the modern electronic limit order book ([Bouchaud et al., 2002](#); [Cont et al., 2010](#); [Gould et al., 2013](#)) and asymmetries in information flow between demand and supply sides ([Korajczyk and Murphy, 2019](#)). Closer to our concerns, the Autoregressive Conditional Duration framework of [Engle and Russell \(1998\)](#); [Engle \(2000\)](#) and the VPIN measure of [Easley et al. \(2011\)](#) link inter-event timing to information arrival, while [Ané and Geman \(2000\)](#) provide direct support for the activity-time hypothesis. Studies of book responses to information events ([Brandt and Kavajecz, 2004](#); [Riordan and Storkenmaier, 2012](#); [Moinas et al., 2020](#); [Bonart and Gould, 2018](#)) and recent machine-learning approaches ([Ntakaris et al., 2018](#); [Zhang et al., 2019](#); [Briola et al., 2024](#)) document rich quote dynamics.

The activity-time hypothesis traces to [Clark \(1973\)](#), building on [Mandelbrot and Taylor \(1967a\)](#). The Lévy-process framework provides the natural mathematical language: [Madan and Seneta \(1990\)](#) and [Carr et al. \(1998\)](#) introduce the Variance Gamma model, [Barndorff-Nielsen \(1997, 1998\)](#) develop the NIG process, which is the identified reduction of our model and whose closed-form density and moments we use in estimation, and [Carr et al.](#)

(2002) and Cont and Tankov (2004) argue that essentially all option-pricing models admit time-changed Lévy representations. Related stochastic-volatility and jump-diffusion work (Heston, 1993; Duffie et al., 2000; Huang and Wu, 2004; Pan, 2002) motivates richer clock structures; more recent contributions document information in the clock beyond realized variance (Li and Zinna, 2018; Carr and Wu, 2009) and provide foundational time-change results (Geman et al., 2001). Multivariate subordinated and tempered-stable frameworks (Semeraro, 2022; Ballotta et al., 2024; Rachev et al., 2011; Bianchi et al., 2011) and the multiple-subordinator approach of Shirvani et al. (2020), extended by Shirvani et al. (2021) and Shirvani et al. (2024), use composed or multiple subordinators to generate heavy tails. Most of this literature estimates a single market-wide clock. Our identification result is consistent with that choice: a composed unit-mean clock is not identified from return cumulants up to fourth order, so single subordination is the specification that four cumulants can identify. Our departure is to model the two sides of the quote jointly, with a common price component and a stationary spread. A further feature of the analysis concerns the law of the clock. This literature typically takes the clock to be a Lévy subordinator, which fixes the rate at which returns approach Gaussianity under aggregation; Section 6.2 measures that rate on tick data and finds it to be roughly half of what any member of the class implies. The non-identification result does not depend on the clock’s law, since it concerns what cumulants can reveal, but the aggregation evidence may be informative about the class of clocks for which such models are appropriate.

Finally, the uncertainty-measurement literature spans options-implied indices (Whaley, 2000; Bollerslev et al., 2009; Bekaert et al., 2013; Drechsler and Yaron, 2011); macroeconomic and policy-based measures (Jurado et al., 2015; Baker et al., 2016; Bloom, 2009; Caldara and Iacoviello, 2022); structural and option-implied higher-moment estimates (Ludvigson et al., 2021; Bali et al., 2021; Kelly and Jiang, 2014); and disagreement-based measures (Buraschi et al., 2014; Banerjee and Kremer, 2010; Carlin et al., 2014). These proxies are derived from transaction prices, option prices, aggregate news indices, or survey data. The SUI is constructed from quotes, and it separates the magnitude of quote uncertainty from the irregularity of its timing.

Positioning for financial stability. The paper’s central question is whether the timing of quote revisions adds information about short-lived episodes of market liquidity stress beyond the magnitude of price movements. The microstructure literature provides the economic setting for quote formation, while the time-change literature provides a tractable way to represent irregular activity. We use that framework to construct a measure of clock dispersion, relate it to observable mid-quote revisions, and examine its behavior around scheduled policy announcements. The identification analysis and the tests of aggregation are supporting evidence about the measurement framework. They also impose important limits on interpretation. In particular, the evidence does not establish that clock dispersion forecasts future liquidity conditions or measures information arrival directly. Its narrower role in this paper is descriptive: it captures how concentrated and irregular price revisions are within an episode of market stress.

3 Model and Measurement

This section develops a framework for the evolution of quotes and the uncertainty that emerges prior to trade execution, examines which of its parameters are identified from returns, and states the estimators.

3.1 Economic Environment

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space with filtration $\{\mathcal{F}_t\}_{t \geq 0}$ representing public information. We introduce stochastic business time $T(t)$, an increasing, right-continuous process with $T(0) = 0$, whose increments measure the amount of market activity in each calendar interval. Under the normalization used below its mean rate is one, and its dispersion measures how irregularly activity arrives.

The economy contains a risk-free asset B_t and a risky asset. The risk-free asset evolves as $dB_t = rB_t dt$, where $r \geq 0$ is constant. Let S_t^b and S_t^a denote the best bid and ask prices. The transaction price S_t is realized only upon a trade:

$$S_t^b \leq S_t \leq S_t^a. \quad (1)$$

The bid-ask spread $\Delta_t = S_t^a - S_t^b$ captures transaction costs, inventory compensation, adverse-selection risk, and disagreement between the two sides. Agents form expectations about fundamental value based on $\mathcal{F}_t$ and private signals; their reservation prices aggregate through the limit order book into observed bid and ask quotes.

Bid and ask prices may evolve under side-specific stochastic business clocks $T_b(t)$ and $T_a(t)$, which allows the two sides to differ in the irregularity with which they incorporate information. Fluctuations in uncertainty can then arise from changes in the *magnitude* of price movements and from changes in the *dispersion* of business time.

3.2 Dynamic Asset Prices

Let $X_t^b = \log S_t^b$ and $X_t^a = \log S_t^a$. Both processes evolve in stochastic business time, motivated by the activity-time hypothesis (Clark, 1973; Carr et al., 2002).

Baseline time-changed dynamics. For each side $i \in \{b, a\}$:

$$X_t^i = X_0^i + \mu_i t + \sigma_i B_{T_i(t)}, \quad (2)$$

where B is a standard Brownian motion evaluated at stochastic time $T_i(t)$. Increments are conditionally Gaussian in business time but non-Gaussian in calendar time (Mandelbrot and Taylor, 1967b; Duffie et al., 2000).

Lévy subordination and heavy-tailed dynamics. To capture the skewness and fat tails of high-frequency quote returns, define the composite side-specific business clock $\Theta_i(t) = T_i(U_i(t))$, where U_i and T_i are independent Lévy subordinators. The log-price process is

$$X_t^i = X_0^i + \mu_i t + \rho_i \Theta_i(t) + \sigma_i B_{\Theta_i(t)}, \quad \Theta_i(t) = T_i(U_i(t)). \quad (3)$$

The drift μ_i accrues in calendar time and ρ_i in business time. Conditional on the clock, $X_t^i - X_0^i - \mu_i t \mid \Theta_i(t) \sim N(\rho_i \Theta_i(t), \sigma_i^2 \Theta_i(t))$, so X^i is a normal mean–variance mixture. We call ρ_i the *business-time drift*. Three of its properties matter for what follows. First, ρ_i is a within-side parameter; it is not a measure of dependence between the bid and ask sides, which is specified separately below. Second, ρ_i generates return skewness: the third cumulant in (6) has the sign of ρ_i . Third, when $\rho_i = 0$, X^i is a Gaussian variance mixture with zero skewness, and the stochastic clock enters the observable moments only through kurtosis. Double subordination is a standard device for generating heavy tails and clustering (Cont, 2001; Rachev and Mittnik, 2000; Kim et al., 2010; Rachev et al., 2011).

Scope of the Lévy assumption. Taking T and U to be Lévy subordinators makes the derivations below tractable: the clock has independent, stationary increments, so the cumulants of X are linear in the horizon, the composed clock has the closed form (4), and the estimated model is exactly normal inverse Gaussian. The assumption also has an observable implication, which we examine in Section 6.2. Linearity of the cumulants in the horizon fixes the rate at which returns approach Gaussianity under aggregation: excess kurtosis falls as $1/\tau$ for returns sampled every τ units. The tick data we use are not fully consistent with that rate, and this bears on how the results should be read. The identification results concern what return cumulants can reveal about a clock whatever law the clock has, and they do not depend on the assumption. The estimators recover the four cumulants of the return distribution and remain valid as cumulant-matching procedures. The interpretation of v as the variance of a Lévy clock increment is less secure, however, and because $v(\tau)$ need not scale as the model implies, a value estimated at one sampling frequency may not be comparable with one estimated at another. We therefore report the sampling interval with every estimate, and Section 6.3 examines what the fitted clock appears to track.

Cumulants. Let T and U be inverse-Gaussian subordinators normalized to unit mean, $\mathbb{E}[T(1)] = \mathbb{E}[U(1)] = 1$, with $\text{Var}(T(1)) = 1/\lambda_T$ and $\text{Var}(U(1)) = 1/\lambda_U$. Their Laplace exponents are $\phi_j(z) = \lambda_j(\sqrt{1 + 2z/\lambda_j} - 1)$, $j \in \{T, U\}$, and composition gives $\mathbb{E}[e^{-z\Theta(t)}] = \exp\{-t\phi_U(\phi_T(z))\}$. Differentiating $\phi_U \circ \phi_T$ at zero yields the cumulants of $\Theta(1)$:

$$\begin{aligned} \kappa_1(\Theta) &= 1, & \kappa_2(\Theta) &= \frac{1}{\lambda_T} + \frac{1}{\lambda_U} \equiv v_\Theta, \\ \kappa_3(\Theta) &= 3 \left(\frac{1}{\lambda_T^2} + \frac{1}{\lambda_T \lambda_U} + \frac{1}{\lambda_U^2} \right), & \kappa_4(\Theta) &= \frac{15}{\lambda_T^3} + \frac{15}{\lambda_T^2 \lambda_U} + \frac{18}{\lambda_T \lambda_U^2} + \frac{15}{\lambda_U^3}. \end{aligned} \quad (4)$$

Because the cumulant generating function of X is $K_X(s) = \mu s + K_\Theta(\rho s + \frac{1}{2}\sigma^2 s^2)$, the cumulants of the log-price increment per unit of calendar time are

$$\begin{aligned} \kappa_1(X) &= \mu + \rho, \\ \kappa_2(X) &= \rho^2 v_\Theta + \sigma^2, \end{aligned} \quad (5)$$

$$\begin{aligned} \kappa_3(X) &= \rho^3 \kappa_3(\Theta) + 3\rho \sigma^2 v_\Theta, \\ \kappa_4(X) &= \rho^4 \kappa_4(\Theta) + 6\rho^2 \sigma^2 \kappa_3(\Theta) + 3\sigma^4 v_\Theta. \end{aligned} \quad (6)$$

The two subordinators enter the variance additively and symmetrically; there is no interaction term. Double subordination shows up only in the third and fourth cumulants.

Proposition 1 (Non-identification of the composed clock). *Write $v_\Theta = 1/\lambda_T + 1/\lambda_U$ and $\omega = (1/\lambda_T)/v_\Theta \in (0, 1)$. Then*

- (a) $\kappa_2(\Theta) = v_\Theta$ and $\kappa_3(\Theta) = 3v_\Theta^2(1 - \omega + \omega^2)$ are invariant to exchanging λ_T and λ_U ($\omega \mapsto 1 - \omega$);
- (b) the exchange changes the fourth cumulant by exactly $\kappa_4(\Theta; \omega) - \kappa_4(\Theta; 1 - \omega) = 3\omega(1 - \omega)(1 - 2\omega)v_\Theta^3$, whose modulus is at most $\frac{\sqrt{3}}{6}v_\Theta^3$;
- (c) for X , the exchange leaves $\kappa_1(X)$, $\kappa_2(X)$ and $\kappa_3(X)$ unchanged and moves $\kappa_4(X)$ by ρ^4 times the amount in (b); at $\rho = 0$ the clock enters the first four cumulants of X only through v_Θ ;
- (d) the first four cumulants of X do not identify $(\mu, \rho, \sigma, \lambda_T, \lambda_U)$; the restriction $U(t) = t$ (single subordination) yields an exactly identified model (Lemma 1).

Proof. Substitute $1/\lambda_T = \omega v_\Theta$ and $1/\lambda_U = (1 - \omega)v_\Theta$ into (4). Then $\kappa_2(\Theta) = v_\Theta$ and $\kappa_3(\Theta) = 3v_\Theta^2(\omega^2 + \omega(1 - \omega) + (1 - \omega)^2) = 3v_\Theta^2(1 - \omega + \omega^2)$, which is symmetric about $\omega = 1/2$. This proves (a). Next, $\kappa_4(\Theta) = v_\Theta^3 q(\omega)$ with $q(\omega) = 15\omega^3 + 15\omega^2(1 - \omega) + 18\omega(1 - \omega)^2 + 15(1 - \omega)^3$, and $q(\omega) - q(1 - \omega) = \omega(1 - \omega)[15\omega + 18(1 - \omega) - 15(1 - \omega) - 18\omega] = 3\omega(1 - \omega)(1 - 2\omega)$. On $[0, 1]$ the modulus of $g(\omega) = \omega(1 - \omega)(1 - 2\omega)$ is maximized where $6\omega^2 - 6\omega + 1 = 0$, i.e. $\omega = (3 \mp \sqrt{3})/6$, where $\omega(1 - \omega) = 1/6$ and $|1 - 2\omega| = \sqrt{3}/3$, so $|3g| \leq \sqrt{3}/6$. This proves (b). Part (c) follows from (5)–(6): $\kappa_1(X)$ does not involve the clock, $\kappa_2(X)$ and $\kappa_3(X)$ involve it only through v_Θ and $\kappa_3(\Theta)$, and $\kappa_4(\Theta)$ enters $\kappa_4(X)$ with coefficient ρ^4 ; at $\rho = 0$, $\kappa_3(X) = 0$ and $\kappa_4(X) = 3\sigma^4 v_\Theta$. For (d), the map from the five parameters to $(\kappa_1, \dots, \kappa_4)$ has fibres of dimension at least one; with $U(t) = t$ the recursion of Lemma 1 inverts it uniquely on its range. $\square$

Maximizing $|q(\omega) - q(1 - \omega)|/q(1 - \omega)$ over $[0, 1]$ shows that the exchange changes $\kappa_4(\Theta)$ by at most 2.96% (at $\omega \approx 0.25$ or 0.75). Which subordinator carries the clock dispersion is therefore effectively unidentified, and so is any quantity, such as an index built on λ_T alone, that depends on it. The estimated model in this paper is the identified reduction, $U(t) = t$; double subordination remains the general framework that motivates it.

Joint specification of the two sides. Treating the bid and ask as independent subordinated processes leaves their joint law unspecified and is not consistent with (1): two independent processes cross with positive probability. At the parameter values estimated in Section 5, independent sides would cross within 30 minutes of the open on a median 62% (2008) and 18% (2025–26) of simulated paths, and within the day on 90% and 78%, whereas the quotes in our data never cross. The data are also difficult to reconcile with independence directly: one-minute bid and ask returns have a correlation of 0.99, and the log spread is strongly mean reverting rather than a random walk (Section 4). We therefore write the two sides as a common price component and a positive spread,

$$X_t^b = X_t^m - \frac{1}{2}s_t, \quad X_t^a = X_t^m + \frac{1}{2}s_t, \quad s_t = e^{\ell_t}, \quad (7)$$

where the mid-quote X^m follows (3) and the log relative spread ℓ is an Ornstein–Uhlenbeck process, possibly run in the same business time, whose innovations have correlation ϑ with those of X^m . The quotes cannot cross because $s_t > 0$, bid–ask asymmetry becomes a statement about the spread process, and ϑ , together with the shared clock, provides an explicit cross-side dependence parameter. In this paper we estimate the marginal single-subordination model separately for the bid, the ask, and the mid-quote. Estimating (7) would require an instrument whose spread has room to move, and we leave it to future work (Section 7).

3.3 Identification and Estimation

Lemma 1 (Moment identification under single subordination). *Let $U(t) = t$, so that $\Theta = T$, $v \equiv v_\Theta = 1/\lambda_T$, $\kappa_3(\Theta) = 3v^2$ and $\kappa_4(\Theta) = 15v^3$. Write $g_1 = \kappa_3(X)/\kappa_2(X)^{3/2}$ and $g_2 = \kappa_4(X)/\kappa_2(X)^2$. Then*

$$\kappa_3(X) = 3\rho v \kappa_2(X), \quad \kappa_4(X) = 3v \kappa_2(X)^2 + \frac{4}{3} \kappa_3(X)^2 / \kappa_2(X),$$

and the parameters are recovered recursively:

$$v = \frac{3g_2 - 4g_1^2}{9}, \quad \rho = \frac{\kappa_3(X)}{3\kappa_2(X)v}, \quad \sigma^2 = \kappa_2(X) - \rho^2 v, \quad \mu = \kappa_1(X) - \rho. \quad (8)$$

An admissible solution ($v > 0$, $\sigma^2 > 0$) exists if and only if $3g_2 > 5g_1^2$.

Proof. Insert $\kappa_n(\Theta)$ into (5)–(6). With $a = \rho^2 v$, $\kappa_3(X) = 3\rho v(\rho^2 v + \sigma^2)$ and $\kappa_4(X) = 3v(5a^2 + 6a\sigma^2 + \sigma^4) = 3v\kappa_2(X)(\kappa_2(X) + 4a)$. Eliminating $a = \kappa_3(X)^2/(9\kappa_2(X)^2 v)$ gives the second identity; solving it for v gives (8). Finally, $\sigma^2 > 0 \iff v > g_1^2/9 \iff 3g_2 > 5g_1^2$. $\square$

The identification is recursive:

Moment	Identifies
$\kappa_1(X) = \mu + \rho$	the sum only; μ is separated from ρ through κ_3
$\kappa_2(X) = \rho^2 v + \sigma^2$	total variance
$\kappa_3(X) = 3\rho v \kappa_2(X)$	the product ρv ; its sign is the sign of ρ
$\kappa_4(X)$	v , given $\kappa_2(X)$ and $\kappa_3(X)$

Two consequences follow. First, the five-parameter vector $(\mu, \sigma, \rho, \lambda_T, \lambda_U)$ matched to four moments is underidentified before any small-sample consideration; a moment-matching criterion can then be driven to zero along a continuum of parameter values, so the estimates of λ_T that such a criterion returns are not determined by the moments. Second, as $\rho \rightarrow 0$, $\kappa_3(X) \rightarrow 0$ and $g_2 \rightarrow 3v$: the clock is identified from excess kurtosis alone, while its contribution to the variance, $\rho^2 v$, vanishes. A fourth cumulant is also a noisy statistic in small samples. The moment estimator of excess kurtosis cannot exceed $n - 5 + 1/(n - 1)$, which is 5.1 for $n = 10$, and its standard error under normality is $\sqrt{24/n}$, which is 1.5 at $n = 10$. We therefore estimate on full trading days of one-minute returns ($n = 389$); a Monte Carlo experiment in the replication files indicates that at this sample size the maximum likelihood estimator of $\log v$ has a root-mean-squared error of 0.3–0.6 and is nearly unbiased, whereas at $n = 30$ it exceeds 2.

Characteristic function. Conditioning on the clock,

$$\mathbb{E}[e^{iuX_t}] = \exp\left(iu\mu t - t\phi_U\left(\phi_T\left(\frac{1}{2}\sigma^2 u^2 - iu\rho\right)\right)\right), \quad (9)$$

with the principal branch of the square root. The composition of two inverse-Gaussian subordinators is not inverse Gaussian, so under double subordination X is not NIG. Under single subordination ($U(t) = t$), (9) reduces to the NIG characteristic function, and the density of X_1 is

$$f(x) = \frac{\alpha\delta K_1(\alpha\sqrt{\delta^2 + (x - \mu)^2})}{\pi\sqrt{\delta^2 + (x - \mu)^2}} \exp\{\delta\gamma + \beta(x - \mu)\},$$

where $\beta = \rho/\sigma^2$, $\gamma = \sqrt{\lambda_T}/\sigma$, $\delta = \sigma\sqrt{\lambda_T}$ and $\alpha = \sqrt{\gamma^2 + \beta^2}$ (so that $\delta\gamma = \lambda_T$). The NIG specification is therefore exact in the estimated model.

Normalization. The normalization $\mathbb{E}[T(1)] = 1$ fixes the mean rate of business time at one unit per unit of calendar time. It therefore identifies the *dispersion* of the clock, not its speed: $\lambda_T = 1/v$ governs $\text{Var}(dT)$, and a larger λ_T means a smoother, more regular clock. Neither v nor λ_T measures the rate at which information arrives. If an intensity object is wanted, the direct route is the observed number of quote revisions per unit of time, which is measurable from tick data and is measured in Section 6.3; the one-minute snapshots used here change in at least 98.7% of minutes, so a per-minute revision count is censored at one.

Estimators. We estimate the single-subordination model by maximum likelihood on each trading day's one-minute returns, after dividing returns by the square root of an intraday variance profile estimated on control days (Andersen and Bollerslev, 1997) (Section 4), and we report the closed-form estimator (8), which is the exactly identified GMM estimator, as a check. The likelihood is maximized over $(\mu, \rho, \log \sigma, \log v)$ on standardized returns with bounds $\log \sigma \in [\log 10^{-3}, \log 10]$ and $\log v \in [\log 10^{-6}, \log 10^4]$, from up to seven starting values (the moment estimate and a grid in ρ and $\log v$). Confidence intervals are percentile intervals from a parametric bootstrap that re-estimates the model on samples simulated from the fitted parameters ($B = 199$ for the likelihood estimator and $B = 999$ for the moment estimator); we report the share of bootstrap samples for which the moment estimator is inadmissible. We do not use characteristic-function matching.

3.4 Model-Implied Variance Decomposition

In a time-changed model, the variance of returns combines a diffusion component and a component that exists only when the business-time drift is non-zero. From (5), per unit of calendar time,

$$\text{Var}(X) = \rho^2 v_\Theta + \sigma^2, \quad v_\Theta = \frac{1}{\lambda_T} + \frac{1}{\lambda_U}, \quad (10)$$

with $v_\Theta = 1/\lambda_T$ in the estimated model. Two periods with identical volatility can differ in the dispersion of business time; the variance registers that difference only through $\rho^2 v_\Theta$, whereas the fourth cumulant (6) registers it directly.

3.5 Definition of the SUI

We define the index as the model-implied return volatility per unit of time:

$$\text{SUI}_t = \sqrt{\rho_t^2 v_{\Theta,t} + \sigma_t^2}, \quad s_t^D = \frac{\sigma_t^2}{\text{SUI}_t^2}, \quad s_t^C = \frac{\rho_t^2 v_{\Theta,t}}{\text{SUI}_t^2}, \quad s_t^D + s_t^C = 1. \quad (11)$$

The index includes the diffusion variance and the clock variance scaled by the business-time drift; the parameter ρ enters only through the clock term, which is where the model places it. SUI_t is a return standard deviation per unit of time and is annualized exactly like realized volatility, by $\sqrt{390 \times 252}$ for one-minute returns. The diffusion share s_t^D and the clock share s_t^C are unit-free and invariant to the time unit: under $t \mapsto ct$, $\rho \mapsto c\rho$, $v_{\Theta} \mapsto v_{\Theta}/c$ and $\sigma \mapsto \sqrt{c}\sigma$.

Remark 1. Under the exactly identified estimator (8), $\widehat{\text{SUI}}^2 = \hat{\kappa}_2(X)$, so the estimated index equals the sample standard deviation of returns in the estimation window. The information the model adds is in the decomposition. By (8), $s^C = g_1^2/(3g_2 - 4g_1^2)$, which is identified by skewness and is zero when $g_1 = 0$. The clock dispersion v , identified by kurtosis, measures irregularity of business time that the variance share cannot register when $\rho = 0$. We therefore report the pair (SUI_t, v_t) : the first for the magnitude of pre-trade uncertainty, the second for the irregularity of its timing.

3.6 Bid–Ask Construction

For each trading day the model is estimated separately on bid, ask, and mid-quote returns, giving

$$\text{SUI}_t^b = \sqrt{(\rho_t^b)^2 v_t^b + (\sigma_t^b)^2}, \quad \text{SUI}_t^a = \sqrt{(\rho_t^a)^2 v_t^a + (\sigma_t^a)^2}. \quad (12)$$

Differences between SUI_t^b and SUI_t^a are differences between the two fitted marginal processes; under (7) they are statements about the dynamics of the spread.

3.7 Exact Decomposition of the Bid–Ask Uncertainty Gap

Proposition 2 (Exact decomposition of the bid–ask uncertainty gap). *For $i \in \{b, a\}$ let $D_i = \sigma_i^2$, $C_i = \rho_i^2 v_i$ and $S_i = D_i + C_i = \text{SUI}_i^2$. Let $\text{LM}(p, q) = (p - q)/(\ln p - \ln q)$ for $p \neq q$ and $\text{LM}(p, p) = p$ (the logarithmic mean). Then, exactly,*

$$\Delta L \equiv \ln \text{SUI}_b - \ln \text{SUI}_a = \underbrace{\frac{D_b - D_a}{2 \text{LM}(S_b, S_a)}}_{\text{diffusion channel}} + \underbrace{\frac{C_b - C_a}{2 \text{LM}(S_b, S_a)}}_{\text{clock channel}}. \quad (13)$$

Proof. $\ln \text{SUI}_b - \ln \text{SUI}_a = \frac{1}{2}(\ln S_b - \ln S_a) = (S_b - S_a)/(2 \text{LM}(S_b, S_a))$ by the definition of LM; substitute $S_b - S_a = (D_b - D_a) + (C_b - C_a)$. $\square$

The decomposition requires no reference point and no approximation, and the two channels sum to the total identically.

Two caveats accompany it. The gap is well defined only once the two sides are modeled jointly: under the independent-sides specification the quotes cross, which the data never

do, so $\ln \text{SUI}_b - \ln \text{SUI}_a$ inherits whatever the dependence assumption supplies (Section 3, equation (7)). And it is an object about *quotes*, not about liquidity provision: an index has no dealer, and, as Section 6.1 indicates, even in market quotes with depth the two sides of a penny-tick instrument move together with a correlation of 0.9994 at one-minute sampling. Proposition 2 is therefore exact arithmetic whose empirical content appears modest on both of the instruments used here.

3.8 Interpretation of the Clock Parameters

Remark 2 (Clock increments). Under the normalization $\mathbb{E}[T(1)] = 1$, increments satisfy $T(t+h) - T(t) \sim \text{IG}(h, \lambda_T h^2)$, with mean h and variance h/λ_T . The parameter λ_T therefore governs the dispersion of business time, not its speed: a larger λ_T means a more regular clock, and $v = 1/\lambda_T$ rises when activity arrives in bursts. An economic reading of a high v is that activity, and possibly the incorporation of information, is concentrated in a few intervals; it is an interpretation, not a measurement of information flow.

Remark 3 (Conditional drift). From (3), $\mathbb{E}[dX_t \mid d\Theta(t)] = \mu dt + \rho d\Theta(t)$.

3.9 Estimation Windows

The index is estimated on full trading days (389 one-minute returns) rather than on short rolling windows, for the reasons given in Section 3.3. Event-level comparisons use model-free statistics in windows around each documented announcement (Section 4), since 30 or 60 one-minute returns appear too few to estimate the clock: in the replication files the 95% bootstrap intervals for v from 30 returns span several orders of magnitude, and those for the clock share are close to $[0, 1]$. Estimates are not winsorized.

4 Data and Empirical Design

Instrument and source. Our index-level quotes are the best bid and best ask of the S&P 500 index (SPX), recorded at a one-minute frequency from Bloomberg’s $< GP >$ function. We do not observe how these index-level quotes are constructed from constituent quotes, or the update frequency of the underlying feed, and we therefore treat them as an index-level quote series. An index has no dealer, depth, or inventory, so these series measure index-level quote dispersion, not liquidity provision, and every statement in this paper about these quotes refers to that object. Statements about liquidity provision require genuine two-sided quotes; Section 6 accordingly examines the framework on SPY National Best Bid and Offer (NBBO) quotes and finds that, at one-minute sampling, the two sides of the NBBO are, if anything, closer together than the two sides of the index. The 2008 one-minute quotes were retrieved from the same source, and we do not have independent information on whether they were recorded contemporaneously or reconstructed; the 2008 evidence is therefore best read as illustrative.

Sample. The data comprise five files of consecutive trading days: 14–20 March 2008 (Bear Stearns, 5 days), 12–19 September 2008 (Lehman Brothers, 6 days), 15 January–14 February 2025 (tariff news, 22 days), 8–15 December 2025 (FOMC, 6 days) and 23–30

January 2026 (FOMC, 6 days), for a total of 45 trading days: 11 in 2008 and 34 in 2025–26. We use the regular session from 09:30 to 15:59 ET, which gives 389 one-minute returns per day, computed within the day. The 16:00 stamp is excluded: in 2025–26 the index quote widens mechanically at the close, to a median quoted spread of 174 basis points against about 5 basis points during the session, and that single minute would dominate the higher moments of the daily return distribution. We also remove 800 post-close rows (to 16:19) and 820 exact duplicate rows (every row of 30–31 January 2025 appears twice in the source file). No session minute is missing.

Quote quality. Table 1 reports diagnostics by file. The series do not appear stale: the share of zero one-minute returns never exceeds 1.03% (2008) or 1.29% (2025–26) on any day or side, and no quote is unchanged for more than one minute. No quote is crossed. The quoted spread has a median of 6.8 basis points in 2008 and 5.0 in 2025–26 and follows a pronounced intraday pattern, falling from 16–19 basis points at the open to 3–7 basis points in the afternoon (Figure 1). The bid and ask sides move together: the median daily correlation of one-minute bid and ask returns is 0.994 (2008) and 0.990 (2025–26), and each side shares about 99.5% of its return variance with the mid-quote (Figure 2). At one-minute resolution the two sides are close to the same price process, to within about half a percent of variance, which is why we estimate the model on the mid-quote as well as on each side. Mid-quote returns in 2008 show the small positive first-order autocorrelation typical of an index with non-synchronous constituent prices (median 0.04).

Table 1. One-Minute SPX Quote Series: Coverage and Diagnostics

Sample	Days	Vol. (%)	Spread (bp)		Zero ret. (%)	Corr. (r^b, r^a)	Common (%)	AC(1)	Exc. kurt.
	(control)		Intraday	16:00				mid	
Bear Stearns, March 2008	5 (3)	26.6	5.1	10	1.03	0.997	99.9	0.005	1.72
Lehman, September 2008	6 (4)	36.7	8.9	14	1.03	0.994	99.7	0.081	0.65
Tariff period, Jan.–Feb. 2025	22 (20)	8.0	5.2	194	1.29	0.990	99.5	0.032	2.68
FOMC, December 2025	6 (5)	9.9	4.3	111	1.03	0.991	99.5	0.011	2.50
FOMC, January 2026	6 (5)	7.7	4.4	109	1.29	0.987	99.3	−0.005	3.75

Notes: One-minute SPX bid and ask quotes, 09:30–15:59 ET (389 returns per day); entries are medians across days unless stated. Days: trading days, with control days in parentheses; control days exclude the event days and the other FOMC statement days in the files (18 March 2008, 16 September 2008, 29 January 2025). Vol.: annualized standard deviation of one-minute mid-quote log returns. Spread: quoted spread $(S^a - S^b)/S^m$, intraday median and at the 16:00 stamp, which is excluded from the analysis. Zero ret.: largest share of zero one-minute returns on any day and side. Corr.: correlation of bid and ask returns. Common: R^2 of each side’s return on the mid-quote return, averaged over the two sides. AC(1) and Exc. kurt.: first-order autocorrelation and excess kurtosis of mid-quote returns.

Event timing and calendars. Table 2 reports the timing of each event, checked against public sources. Both 2008 announcements preceded the open: Bear Stearns announced its secured facility with JPMorgan Chase and the Federal Reserve Bank of New York by press release on the morning of 14 March 2008, and Lehman Brothers filed its Chapter 11 petition at 01:45 on 15 September. Neither has an intraday announcement time, so both are analyzed as full trading days. The two FOMC statements were released at 14:00 ET and followed by the Chair’s press conference at 14:30. The documented intraday tariff event in




Figure 1. Intraday Profiles of Volatility and the Quoted Spread. *Notes:* Control days only, 09:30–15:59 ET; 9-minute centered moving averages, shown separately for the 2008 and 2025–26 samples. Left: annualized volatility of one-minute mid-quote returns (logarithmic scale). Right: median quoted spread in basis points.




Figure 2. Co-Movement of the Bid and Ask Sides. *Notes:* One marker per trading day. Left: correlation of one-minute bid and ask returns. Right: share of each side’s return variance explained by the mid-quote return (R^2); circles denote the bid and triangles the ask.

the 2025 file is Mexico’s one-month tariff pause on 3 February 2025, announced by Mexico’s president on X at 10:21 ET; the index quote reaches its low at 10:20 and is 0.85% higher five minutes later. Canada’s pause came after the close. Tariff remarks on the evening of 20 January (a market holiday) and after the close on 21 January fell outside the session, so those days are not treated as intraday events. For the panel of Section 6 we compiled from public sources the dates and announcement times of all 146 FOMC statements from 2010 to September 2026, the dates of 200 CPI and 200 Employment Situation releases (08:30 ET) over the same period, and the high-frequency monetary-policy surprises of [Jarociński and Karadi \(2020\)](#), which are available through March 2026.

Table 2. Event Timing and Windows

Event	Date	Announcement (ET)	Window used	Source
Bear Stearns facility	14 Mar 2008	Morning, pre-open	Full day	SEC DEFM14A; Federal Reserve release
Lehman Chapter 11 filing	15 Sep 2008	01:45, pre-open	Full day	Federal Reserve History; bankruptcy record
FOMC statement (25 bp cut)	10 Dec 2025	14:00; press conf. 14:30	14:00 $\pm$ 60 min	Federal Reserve statement
FOMC statement (hold)	28 Jan 2026	14:00; press conf. 14:30	14:00 $\pm$ 60 min	Federal Reserve statement and press-conference page
Tariff remarks	21 Jan 2025	None in session	Not used	CNBC, 21 January 2025
Tariff: Mexico pause	3 Feb 2025	10:21:06	10:20 $\pm$ 30 min	Sheinbaum post on X; NBC

Notes: Announcement times are taken from public sources; full citations are in the replication files. Both 2008 announcements preceded the 09:30 open, so no intraday pre/post window exists and they are analyzed as full trading days. The Mexico tariff pause was announced by Mexico’s president on X at 10:21:06 ET (the first press report followed at 10:41); the window is centered on the 10:20 stamp, the last quote before the post. In the file the index quote reaches its low at 10:20 and is +0.85% higher by 10:25. Canada’s pause was announced after the close (16:36 ET). Tariff remarks on the evening of 20 January (a market holiday) and after the close on 21 January fell outside the session.

The time-of-day confound. Intraday volatility and quoted spreads follow a pronounced U-shape (Figure 1). Any comparison of a window before an event with the window after it therefore mixes the event with the time of day. In the case studies the event types are collinear with clock time: both 2008 events and the first tariff window fall in the morning, and both FOMC events at 14:00. Figure 3 shows how often volatility falls across each clock time on days *without* an event. In the morning it falls on 60–86% of control days in both periods, so morning declines are consistent with the ordinary intraday pattern. At 14:00 the two periods differ: in 2008 volatility usually rises across 14:00 (it falls on 14% of control days), whereas in 2025–26 it usually falls (70% of control days for ± 30 -minute windows and 77% for ± 60). After time-of-day adjustment the shares move toward one half (40–60% in 2025–26; the 2008 profile rests on only seven days).




Figure 3. No-Event Benchmark. *Notes:* Share of control days on which the standard deviation of one-minute mid-quote returns is lower in the 30 minutes after a given clock time than in the 30 minutes before it, by sample period. The dashed line marks one half.

Design. We address the confound in three ways. First, every event window is compared with the same clock-time window on the control days of the same period (non-event days in the same files, excluding other FOMC statement days). Second, returns are also divided by the square root of an intraday variance profile estimated on control days, following [Andersen and Bollerslev \(1997\)](#), with each control day left out of its own profile. Third, once the event set is expanded, event-type and time-of-day fixed effects make the two separately estimable; this is done on the panel in Section 6.4. The first two are implemented in the case studies. With 30 control days in 2025–26 and 7 in 2008, inference is by randomization: an event’s rank among its controls and a two-sided randomization p -value, whose minimum attainable value is $1/31 = 0.032$ in 2025–26 and $1/8 = 0.125$ in 2008. No 2008 event can therefore be statistically distinguished from its controls in these files. With five events in five files, one event per type, and event type collinear with period and volatility regime, the evidence is illustrative case-study evidence; it does not support statements about how event types differ.

Windows. For the FOMC decisions we use ± 60 -minute windows around 14:00, which span the statement and the press conference, and report ± 30 -minute windows as a robustness check (they end as the press conference begins). For the tariff pause we use ± 30 minutes around the 10:20 stamp, the last quote observed before the 10:21 announcement. The 2008 events, which preceded the open, are compared with their control days at the level of the full trading day.

5 Event Evidence

5.1 Full-Day Estimates

Table 3 summarizes the full-day estimates. Three features are worth noting. First, the SUI closely resembles volatility: its median ratio to the sample standard deviation of returns is 0.996–0.997 in every period and on every side, and Figure 4 shows that the daily SUI and realized volatility are nearly the same series. Second, the clock share of variance is

small: its median is 1–3% and its 90th percentile about 8%, and the lower bound of its 95% bootstrap interval exceeds 0.05 on at most one of the 45 days on any side. Once the variance is decomposed, the clock’s contribution to the variance does not appear detectable at the daily level. Third, the clock dispersion v is estimated on every day, with medians of 0.33 minutes (2008) and 0.51 minutes (2025–26), but coarsely: the ratio of the upper to the lower bound of its 95% interval has a median of about 9 in 2025–26 and 28–61 in 2008. The model fits the moments it targets: across all time-of-day-adjusted series the median ratio of fitted to sample variance is 0.994, and the median fitted excess kurtosis (1.64) is close to the median sample excess kurtosis (1.62). The likelihood estimator converged on every adjusted series, and the moment estimator was admissible on 91% (2008) and 100% (2025–26) of them.

Table 3. Full-Day Estimates of the Single-Subordination Model

Sample	Side	Days	SUI		Clock dispersion v (min)		Clock share (%)		MoM
			Ann. (%)	SUI/SD	Median [p10, p90]	CI ratio	Median (p90)	CI upper	
2008	Bid	11	32.2	0.997	0.33 [0.03, 2.02]	28.0	3.0 (7.9)	35.3	91
	Ask	11	32.7	0.997	0.34 [0.05, 1.78]	27.9	2.2 (8.2)	38.9	91
	Mid-quote	11	32.4	0.997	0.33 [0.04, 1.89]	60.6	1.9 (8.5)	66.7	91
2025–26	Bid	34	8.2	0.996	0.51 [0.24, 2.01]	8.5	1.3 (8.0)	13.5	100
	Ask	34	8.1	0.997	0.50 [0.24, 2.04]	9.2	0.9 (7.2)	13.3	100
	Mid-quote	34	8.1	0.997	0.51 [0.24, 2.04]	8.9	1.1 (7.5)	13.1	100

Notes: Single-subordination (NIG) model estimated by maximum likelihood on each trading day’s 389 one-minute returns after time-of-day adjustment; entries are medians across days. SUI is $\sqrt{\rho^2 v + \sigma^2}$, annualized, and SUI/SD its ratio to the sample standard deviation. $v = \text{Var}(T(1))$ is the clock dispersion in minutes; CI ratio is the ratio of the upper to the lower bound of its 95% parametric-bootstrap interval ($B = 199$). Clock share is $\rho^2 v / \text{SUI}^2$, with its 90th percentile in parentheses, and CI upper the upper bound of its interval. MoM admiss.: share of days on which the closed-form moment estimator exists ($3g_2 > 5g_1^2$). The likelihood estimator converged on every adjusted series.

5.2 Event Windows against Control Days

Table 4 and Figure 5 compare each documented intraday announcement with the same window on the control days of its period. All three announcements are followed by an increase in one-minute volatility that ranks at or near the top of the control distribution, and by a widening of the quoted spread that ranks at the top. After the FOMC hold of 28 January 2026, mid-quote volatility rose from 4.8% to 14.6% (annualized) over the ± 60 -minute window, the largest increase among 31 days, and the quoted spread widened by 27 log points, also the largest. Around the FOMC cut of 10 December 2025, volatility rose from 11.8% to 19.6% (29 of 30 controls lower) and the spread widened by 32 log points, the largest in the sample. Around Mexico’s tariff pause, volatility rose from 17.7% to 33.9% (29 of 30 controls lower). These responses run against the intraday pattern: on control days volatility usually falls across 14:00 and 10:20 and the spread narrows. The timing of the FOMC responses may be informative. In the ± 30 -minute windows, which end as the press conference begins, the December volatility response is no longer apparent (21 of 30 controls lower) while the spread widening remains the largest in the sample. For the January decision the largest one-minute moves occur at 14:29–14:30, when the press conference begins.




Figure 4. Daily SUI and Realized Volatility. *Notes:* Mid-quote. SUI is the full-day maximum likelihood estimate on time-of-day-adjusted returns, with 95% parametric-bootstrap intervals (shaded); realized volatility is the root mean square of the day's one-minute returns. Both are annualized and plotted on a logarithmic axis. Open circles mark event days. The four upper panels show the 2008 and FOMC windows; the lower panel shows the tariff period.

The two 2008 events preceded the open, so there is no intraday window for them. With seven control days in that period the smallest attainable two-sided randomization p -value is 0.125, and we do not draw event-level inferences from them; they appear as event days in Figures 4 and 6.

Table 4. Event Windows against the Same Window on Control Days

Event	Window (min)	Mid-quote volatility			Adjusted $\Delta \log$	Quoted spread	
		Pre $\rightarrow$ post (%)	$\Delta \log$	Rank (p)		$\Delta \log$	Rank (p)
FOMC, 10 Dec 2025	± 60	11.8 $\rightarrow$ 19.6	+0.50	29/30 (0.065)	+0.63	+0.32	30/30 (0.032)
FOMC, 10 Dec 2025	± 30 (rob.)	16.0 $\rightarrow$ 16.4	+0.02	21/30 (0.484)	+0.14	+0.26	30/30 (0.032)
FOMC, 28 Jan 2026	± 60	4.8 $\rightarrow$ 14.6	+1.11	30/30 (0.032)	+1.18	+0.27	30/30 (0.032)
FOMC, 28 Jan 2026	± 30 (rob.)	4.8 $\rightarrow$ 17.9	+1.31	30/30 (0.032)	+1.38	+0.14	30/30 (0.032)
Tariff pause, 3 Feb 2025	± 30	17.7 $\rightarrow$ 33.9	+0.65	29/30 (0.065)	+0.75	+0.03	30/30 (0.032)

Notes: Mid-quote one-minute returns in the window before and after the announcement. Volatility is the annualized standard deviation of one-minute returns and $\Delta \log$ the log of the post-event over the pre-event value. Rank is the number of same-period control days (same clock time and window) with a lower value, out of all control days; p is the two-sided randomization p -value $(1 + \#\{|c - \text{med}| \geq |e - \text{med}|\}) / (1 + n)$, whose minimum with 30 controls is 0.032. Adjusted: returns divided by the square root of the control-day intraday variance profile; its ranks equal those of the unadjusted change in every row. (rob.): ± 30 -minute window, which ends as the FOMC press conference begins.




Figure 5. Event Windows against the Same Window on Control Days. *Notes:* Log of the post-event over the pre-event value, for the mid-quote volatility (left), the same volatility adjusted for time of day (center) and the quoted spread (right). Diamonds: the event day; grey dots: the same clock time and window on each control day of the same period.

5.3 FOMC Statement Days and Clock Dispersion

Table 5 and Figure 6 report the full-day clock dispersion on every FOMC statement day in the files. Besides the two decisions studied above, the files contain three statement days that are not events in this paper: 18 March 2008, 16 September 2008 and 29 January 2025. They were excluded from the control days before any estimation. On all five days v is at or near the top of its period's distribution: the two 2008 days rank ninth and eighth of

nine, and the three 2025–26 days rank 33rd, 31st and 30th of 33. Under the null hypothesis that FOMC days are exchangeable with control days within each period, a stratified exact permutation test gives a one-sided p -value of 4.1×10^{-5} for all five days and 0.0027 for the three out-of-sample days alone. The excess appears concentrated in the hour that follows the statement. Dropping the five minutes after 14:00 leaves v at or above the 97th percentile, but dropping the hour 14:00–15:00 returns it to the 77th and 87th percentiles for the two 2025–26 decisions. Daily volatility on those two days is not unusual: the SUI ranks at the 70th and 87th percentiles. (This is not a general property of FOMC days: on 29 January 2025 the SUI was the highest of its period.) By contrast, the tariff-pause day has the highest SUI of 31 days ($p = 0.032$), but its v ranks at the 87th percentile and falls to the median once the hour after 10:20 is removed. Business-time dispersion and volatility thus appear to capture different features of these days. On the two FOMC decision days studied, activity arrives in bursts around the statement and press conference without an unusually volatile day. On the tariff day volatility is high without unusual burstiness. With five statement days this pattern is suggestive only, and Section 6.4 examines it on 166 statement days.

Table 5. FOMC Statement Days and Clock Dispersion

Date	Role	v	Rank	Pct.	Re-estimated without the statement window		
		(min)			v , 14:00–14:05	v , 14:00–15:00	Pct., 14:00–15:00 (%)
18 Mar 2008	Other	4.21	9/9	100	–	–	–
16 Sep 2008	Other	1.89	8/9	100	–	–	–
29 Jan 2025	Other	3.46	31/33	97	–	–	–
10 Dec 2025	Case study	3.57	33/33	100	3.57	1.08	87
28 Jan 2026	Case study	2.25	30/33	97	2.03	0.64	77

Notes: Clock dispersion v (minutes) from the full-day maximum likelihood estimate on time-of-day-adjusted mid-quote returns. Rank is the position of v among the day and the same-period control days (1 = smallest) and Pct. the share of those control days with a smaller v . The last three columns re-estimate the day without the first 5 or 60 minutes after the 14:00 statement (2025–26 days only). The test is a stratified exact permutation test that FOMC days are exchangeable with control days within period (one-sided, sum of within-period ranks): $p = 4.1 \times 10^{-5}$ for all five days (196,416 permutations) and $p = 0.0027$ for the three days that are not case studies, which were excluded from the control days before estimation.

5.4 Bid and Ask

Because the two sides share about 99.5% of their return variance, bid–ask differences in the estimated parameters are small. On 2025–26 control days the bid-side SUI exceeds the ask-side SUI by 0.85% on average ($t = 5.2$ across days), and the model-free ratio of bid to ask volatility points the same way (1.2%, $t = 3.8$); in 2008 there is no systematic gap ($t = -0.7$). The exact decomposition of Proposition 2 holds to machine precision on every day. The 2025–26 event days lie within the range of their controls. The Lehman day has the largest ask-side excess in its period (a gap of -1.7%), which cannot be distinguished from its seven controls. With index quotes, bid–ask asymmetry appears to be a small effect, and it appears smaller still in market quotes: in the SPY NBBO the median daily correlation of one-minute bid and ask returns is 0.9994 (Section 6.1), because the spread is




Figure 6. Daily Clock Dispersion. *Notes:* Maximum likelihood estimate of $v = \text{Var}(T(1))$ (minutes) on time-of-day-adjusted one-minute returns, with 95% parametric-bootstrap intervals, for the bid (circles), the ask (triangles) and the mid-quote (squares), on a logarithmic axis. Shaded: FOMC statement days. The four upper panels show the 2008 and FOMC windows; the lower panel shows the tariff period.

one tick almost always. Identifying a side-specific parameter would require an instrument whose spread has room to move, not merely a real one.

5.5 Comparison with Realized Volatility, Historical Volatility and the Quoted Spread

We compare the index with three benchmarks constructed from the same quotes: realized volatility (RV), the root mean square of the day’s one-minute returns; historical volatility (HV), their standard deviation; and the mean quoted spread. All volatility measures are annualized identically, so the comparison is in levels on a common scale. Table 6 regresses each model object on log RV and the log spread across the 45 days, with a period dummy. The SUI is spanned by realized volatility (R^2 of 0.986 after time-of-day adjustment and 0.999 without). HV is omitted because its daily correlation with RV exceeds 0.99999, and under the exactly identified estimator the SUI equals HV by construction (Remark 1). The clock dispersion is not spanned: log RV and the log spread explain 17% of the daily variation in log v , with a positive loading on RV ($t = 2.4$). In 2025–26 the rank correlation of v with the quoted spread is 0.48, and with the mean absolute change in the spread, a measure of quote activity, 0.53. The clock share is not spanned either ($R^2 = 0.13$), but at the daily level it is mostly estimation noise. With 45 days in five clusters of consecutive days, this comparison is descriptive.

Table 6. Is the Model Spanned by Realized Volatility and the Quoted Spread?

Dependent variable	log RV		log spread		R^2	Days
	Coef.	(t)	Coef.	(t)		
log SUI (time-of-day adjusted)	1.027	(24.3)	0.006	(0.1)	0.986	45
log SUI (raw)	0.985	(79.2)	0.008	(0.6)	0.999	45
log v	1.188	(2.4)	0.049	(0.0)	0.171	45
Clock share	−0.008	(−0.2)	−0.173	(−0.7)	0.129	45

Notes: Daily cross-section of the mid-quote, 45 trading days in five samples; OLS with a period dummy and HC3 t -statistics in parentheses. log RV is the log of annualized realized volatility of one-minute returns and log spread the log of the mean quoted spread (bp). Historical volatility is omitted because its daily correlation with RV is 0.999999; under the exactly identified moment estimator the SUI equals the sample standard deviation (maximum relative difference 7×10^{-13}), so a regression of that SUI on historical volatility has $R^2 = 1$ by construction. Descriptive only: the sample consists of five clusters of consecutive days.

6 Multi-Year NBBO Evidence

The case-study evidence rests on 45 days of index-level quotes. Section 4 noted that statements about liquidity provision, about types of news, and about forecasting require genuine two-sided quotes and a longer sample. This section uses both. We estimate the framework on the SPY National Best Bid and Offer over 779 trading days between January 2007 and December 2025. We also examine the model itself: single subordination has

specific implications for how returns aggregate across sampling frequencies and for what the clock should represent, and tick data allow both to be examined.

6.1 The Panel

We take the complete NBBO for SPY from the TAQ millisecond files and aggregate it, inside the database, into one-second bars carrying the prevailing bid, ask and displayed sizes, the number of NBBO updates and the number of mid-quote changes. The sample is every FOMC statement day, CPI release and Employment Situation release that TAQ covers from 2007 onwards, together with the 30 trading days either side of each case study of Section 5: 779 days, of which 166 are FOMC statement days, 212 CPI releases and 209 Employment Situation releases. Table 7 describes the panel.

Three features are worth noting before any estimation. First, the series are clean: in 18.2 million session seconds not one carries a crossed quote, the NBBO is locked in 0.45% of seconds, and 2.5% of seconds carry no update, in which case the prevailing quote is carried forward exactly rather than imputed. Second, the closing stamp is an ordinary quote here. In the index file the 16:00 spread is 174 basis points against about 5 intraday, and we drop it; in the NBBO it is 0.47 basis points against an intraday median of 0.49, so the session runs to 16:00 and each day contributes 390 one-minute returns. Third, the market these data describe changed over the sample: the median quoted spread falls from 0.78 basis points in 2007–2009 to 0.31 in 2025, displayed depth at the NBBO falls from about 13,800 shares to 1,100, and quote updates rise from 924 a minute to 9,114.

Two limits follow from how the data were drawn. The sample is an event sample, not a random sample of the calendar, so unconditional moments of the panel describe announcement days and the neighborhoods of crises rather than an average trading day; we say so wherever it matters. And WRDS TAQ does not yet carry 2026, so the January 2026 decision studied above has no NBBO counterpart here.

The observation of Section 4 that the two sides of the quote move nearly together carries over to the change of instrument. We estimated the model on the mid-quote as well as on each side because in index quotes the two sides are close to the same process. On market quotes with depth they are closer still: the median daily correlation of one-minute bid and ask returns is 0.9994. In an instrument whose spread is one cent almost always, the two sides cannot move independently at one-minute sampling, so a bid–ask asymmetry in the estimated parameters is not something these data can carry. That is a statement about the tick and the sampling frequency, not about the framework; it is also the reason the remainder of this section works with the mid-quote.

The time-of-day adjustment of Section 4 carries over with one change required by the sample. Because the pull is an event pull, it contains no announcement-free days at all between 2009 and 2023, so a profile estimated from control days alone is unavailable. Each day’s profile is instead built from the 41 pulled days either side of it, using from each contributing day only the minutes outside that day’s own announcement window, and excluding the day being adjusted. The profile is therefore free of announcement contamination but local in event time rather than calendar time: the median profile draws on 82 days spanning about 2.5 calendar years. Since the intraday variance pattern moves slowly, and since the profile enters only as a normalization with mean one, this appears to be a mild approximation; where a result turns on it we report the unadjusted version alongside.

Table 7. The SPY NBBO Panel

Period	Days	Events	Spread	Depth	Updates	Mid-quote	Zero returns (%)		No update
		FOMC/CPI/NFP	(bp)	(000 shares)	per min.	changes per min.	60 s	5 s	(% of seconds)
2007–2009	178	30/23/23	0.779	13.8	924	86	6.7	30.6	5.6
2010–2014	158	43/60/58	0.726	17.6	1452	44	13.1	49.0	4.1
2015–2019	154	41/59/59	0.409	6.1	2323	123	8.7	36.4	5.0
2020–2024	175	44/59/58	0.258	1.0	8397	1345	2.3	9.3	0.1
2025	114	8/11/11	0.314	1.1	9115	1456	1.8	7.6	0.1
All	779	166/212/209	0.488	8.7	2319	171	5.6	25.3	2.5

Notes: SPY National Best Bid and Offer from WRDS TAQ, aggregated on the server into one-second bars labeled by their end stamp, 09:30–16:00 ET. Each cell is the median across the days in the period, except the day and event counts (FOMC statements, CPI releases, Employment Situation releases). Depth is the displayed size at the NBBO, both sides. Zero returns are shares of zero mid-quote returns at the stated sampling interval. No update is the share of seconds in which the NBBO did not change, so the prevailing quote is carried forward exactly rather than imputed. No second in the 779 days carries a crossed quote (0 in total). Six early closes (3 July, the day after Thanksgiving, 24 December) are detected from the collapse in quote activity after 13:00 and truncated there. The sample consists of announcement days plus the 30 trading days on either side of each case study and is not a random draw from the calendar.

Table 8 reports the daily estimates themselves, by period and side, and two features of it are worth having in view before the tests that follow. The moment estimator’s SUI is the sample standard deviation to machine precision on all 2,319 day–side series, as Remark 1 requires, and the likelihood estimator’s is within a tenth of a percent of it at the median; the level of the index carries no information that the sample standard deviation does not. And the clock share of variance is small throughout—a median of about 1% and a 90th percentile below 15% in every period—so the clock’s contribution to the *variance* does not appear to be the object to watch. What varies is v , which is the focus of the rest of this section.

6.2 How Returns Aggregate across Sampling Intervals

Single subordination makes a specific prediction about sampling frequency, which the tick data allow us to examine. X is a Lévy process, so its cumulants are linear in the horizon. For returns sampled over intervals of length τ ,

$$g_1(\tau) = g_1(1) \tau^{-1/2}, \quad g_2(\tau) = g_2(1) \tau^{-1}, \quad (14)$$

and an estimator that reads each return as one unit of time returns $v(\tau) = v(1)/\tau$, while the clock share $\rho^2 v / (\rho^2 v + \sigma^2)$ is invariant. Excess kurtosis must fall as $1/\tau$: on a log–log plot, a line of slope -1 .

Table 9 and Figure 7 report what the data show, at ten sampling intervals from one second to ten minutes. At each τ the cumulants are computed from every offset—the returns beginning at second $0, 1, \dots, \tau - 1$ each form a non-overlapping series, and their k -statistics are averaged—so nothing is discarded and the long-horizon estimates remain usable. The decay is slower than the model implies. The median day’s excess kurtosis is 11.4 at one second and 0.48 at ten minutes: a fall by a factor of 24 across a 600-fold change of horizon, where the model implies a factor of 600. The per-day slope of $\log g_2$ on $\log \tau$ has a median of -0.46 , with a 10th-to-90th-percentile range of -0.74 to -0.21 . Fitted v behaves similarly (median slope -0.50), and the clock share, which the model holds constant, rises monotonically from 0.1% at one second to 7.8% at ten minutes (median

Table 8. Daily Estimates on the SPY NBBO Panel

Period	Side	Days	SUI	Sample	Clock dispersion v			Clock share		MoM	MLE
					(annualized)	s.d.	MLE	95% interval	MoM	Median	p90
2007–2009	Bid	177	0.186	0.188	0.426	[0.09, 0.91]	0.361	0.0115	0.0887	94	100
	Ask	177	0.187	0.188	0.433	[0.08, 0.90]	0.377	0.0109	0.0830	94	100
	Mid-quote	177	0.187	0.188	0.425	[0.07, 0.88]	0.369	0.0114	0.0957	94	100
2010–2014	Bid	157	0.101	0.101	0.666	[0.22, 1.32]	0.553	0.0100	0.0486	97	100
	Ask	157	0.101	0.101	0.672	[0.22, 1.25]	0.559	0.0088	0.0502	97	100
	Mid-quote	157	0.101	0.101	0.663	[0.20, 1.32]	0.557	0.0095	0.0491	97	100
2015–2019	Bid	154	0.086	0.088	0.666	[0.20, 1.28]	0.635	0.0097	0.0443	97	100
	Ask	154	0.087	0.088	0.662	[0.22, 1.28]	0.640	0.0098	0.0461	97	100
	Mid-quote	154	0.086	0.088	0.667	[0.20, 1.23]	0.639	0.0098	0.0448	97	100
2020–2024	Bid	174	0.124	0.125	0.507	[0.12, 1.03]	0.430	0.0138	0.0705	96	99
	Ask	174	0.124	0.125	0.509	[0.13, 1.03]	0.430	0.0131	0.0691	97	99
	Mid-quote	174	0.124	0.125	0.521	[0.13, 1.01]	0.422	0.0128	0.0700	96	100
2025	Bid	111	0.106	0.105	0.385	[0.04, 0.83]	0.328	0.0080	0.1139	94	100
	Ask	111	0.106	0.105	0.389	[0.05, 0.82]	0.326	0.0088	0.1481	94	100
	Mid-quote	111	0.106	0.105	0.389	[0.05, 0.83]	0.327	0.0084	0.1486	94	100

Notes: One-minute bid, ask and mid-quote returns adjusted for time of day with a leave-one-day-out profile; each cell is the median across the days in the period. SUI, the sample standard deviation and the MLE columns use the maximum likelihood fit; MoM is the exactly identified moment estimator, for which the SUI equals the sample standard deviation by construction (to machine precision on all 2,319 day-side series). The likelihood estimator is not bound by that identity and gives a median ratio to the sample standard deviation of 0.999. The interval is the median across days of the 95% parametric-bootstrap percentile interval for v ($B = 199$). MoM admiss.: share of days on which the moment estimator exists ($3g_2 > 5g_1^2$); MLE converged: share of likelihood fits that converge away from the parameter bounds.

slope +0.63).

Two alternative explanations should be considered before reading this as evidence about the clock. The first is the intraday pattern: a deterministic U-shape in variance is itself a mixture and produces excess kurtosis at every horizon. Removing it with the leave-one-day-out profile of Section 4 makes the decay slightly *steeper*, -0.46 against -0.40 on raw returns, and leaves the other results essentially unchanged. The second is price discreteness. SPY trades on a one-cent grid, and at one-second sampling 57% of mid-quote returns are exactly zero, which truncates the tails and could flatten the measured decay. But discreteness is a high-frequency phenomenon: by 60 seconds the zero share is 0.8%, and it is negligible beyond two minutes. Over the range from one minute to ten minutes, where discreteness is unlikely to bind, the slope is -0.46 , close to the slope over the whole range.

A natural reading is the one Cont (2001) gives for this stylized fact: kurtosis that decays more slowly than $1/\tau$ is a signature of a volatility process whose increments are dependent. A single Lévy subordinator has independent increments by construction and is not consistent with it. Two consequences follow, and both shape how we interpret the results below. Whatever the fitted v measures, it is unlikely to be the variance of the increment of a Lévy clock. And because $v(\tau)$ does not scale as the model implies, a value estimated at one sampling frequency should not be carried to another, so every estimate below is reported against a stated frequency.

Table 9. Aggregation across Sampling Frequencies

τ (s)	Excess kurtosis		Clock dispersion v		Clock share	Zero returns
	Data	Model	Data	Model	Data	(%)
1	11.41	97.28	3.768	30.149	0.0012	57.1
2	7.53	48.64	2.481	15.075	0.0016	43.2
5	4.45	19.46	1.463	6.030	0.0023	24.4
10	3.17	9.73	1.041	3.015	0.0038	13.7
15	2.71	6.49	0.868	2.010	0.0048	9.3
30	2.05	3.24	0.656	1.005	0.0073	4.0
60	1.62	1.62	0.502	0.502	0.0114	0.8
120	1.28	0.81	0.379	0.251	0.0220	0.0
300	0.87	0.32	0.230	0.101	0.0517	0.0
600	0.48	0.16	0.089	0.050	0.0780	0.0
Slope on $\log \tau$	-0.458	-1	-0.499	-1	0.631	0

Notes: Mid-quote returns sampled every τ seconds on 773 full-session days, adjusted for time of day. At each τ the cumulants are computed from every offset (the returns starting at second $0, 1, \dots, \tau - 1$ each form a non-overlapping series) and averaged, so no data are discarded. Entries are medians across days. Model columns give the single-subordination benchmark anchored at $\tau = 60$ s: because X is a Lévy process its cumulants are linear in the horizon, so excess kurtosis and v both fall as $1/\tau$ and the clock share is unchanged. The last row is the median across days of the within-day regression of each column on $\log \tau$. Zero returns is the share of zero mid-quote returns, which shows where price discreteness takes over. On raw returns the kurtosis slope is -0.399 rather than -0.458 .




Figure 7. Aggregation across Sampling Frequencies. *Notes:* Mid-quote returns on 773 full-session days, time-of-day adjusted. Solid: the median day's excess kurtosis (left) and fitted clock dispersion v (right) at each sampling interval, with the interquartile range across days shaded. Dashed: what single subordination implies, anchored at $\tau = 60$ seconds (dotted vertical line), namely decay proportional to $1/\tau$. Both axes are logarithmic, so the dashed line has slope -1 .

6.3 Does the Fitted Clock Track Quote Activity?

The clock is the part of the framework that carries an economic interpretation, and it is otherwise latent. Tick data allow a comparison with quote-revision intensity, its observable counterpart. The comparison can be sharper than a correlation. With $E[\Theta(1)] = 1$ and $\text{Var}(\Theta(1)) = v$, a business clock that *is* market activity means the normalized count $N_m/E[N]$ of activity in minute m is the clock increment. That proxy has mean one by construction, so the model predicts a level, not merely co-movement:

$$v = \text{CV}^2(N) = \frac{\text{Var}(N_m)}{E(N_m)^2}. \quad (15)$$

We compare (15) against four candidate clocks: the number of NBBO updates, the number of mid-quote changes, the number of seconds in the minute carrying a non-zero mid-quote return, and the realized variance of the minute’s sixty one-second returns. Counts receive the same leave-one-day-out intraday adjustment as returns, and v is the primary daily estimate from one-minute returns, so the two sides of (15) are measured at the same frequency—which, after Section 6.2, is a requirement rather than a convention. Table 10 and Figure 8 report the results.

Quote activity is highly clustered, which is a precondition for the comparison to be informative. The Fano factor of per-minute NBBO update counts has a median of 772 across days, against one for a Poisson process, and it grows with the bin width at a log–log rate of 0.75; the branching ratio this implies for a linear Hawkes process is 0.98. Quote arrivals are far from evenly spaced.

Co-movement is strong and does not appear to be a trend. The Spearman correlation of $\log v$ with $\log \text{CV}^2$ across the 773 days is 0.78 for NBBO updates and 0.77 for mid-quote changes, and removing year means leaves it at 0.77 and 0.76. Within the day the relation is stronger still: across 8,215 half-hour windows, regressing $\log v$ on $\log \text{CV}^2$ of five-second update counts with day and window-of-day fixed effects gives a coefficient of 1.45 ($t = 28.9$) and an R^2 of 0.62, rising to 0.71 once the window’s own volatility is controlled for. These magnitudes suggest that the fitted clock is related to market activity rather than to an unrelated feature of the returns.

The level restriction is where the candidates separate, and the separation is informative. Raw NBBO update counts are smoother than the fitted clock: their median CV^2 is 0.16 against a median v of 0.53, a ratio of 3.1, with only 19% of days inside a factor of two and an estimated elasticity of 1.59 rather than one. The same holds for the tick clock (ratio 4.7, elasticity 1.57). Mid-quote changes are much closer: a median ratio of 1.24, 61% of days inside a factor of two, an elasticity of 1.14, and a 95% bootstrap interval for v that contains CV^2 on 82% of days. Within-minute realized variance is closer still (ratio 0.77, elasticity 0.996, $p = 0.95$ against one), but it is not an independent comparison: v is identified from the kurtosis of one-minute returns, which is itself driven by variation in within-minute variance, so that row is close to a consistency check on the estimator. The more informative comparison is with mid-quote revisions, which count events in the quote process rather than functions of returns.

One reading is that the fitted clock resembles a *price-revision* clock more than a quote-traffic clock. Most NBBO updates are changes in displayed size or venue that leave the mid-quote where it was; they are numerous, and their count is correspondingly smooth.

The subset that moves the mid-quote is what the fitted clock appears to track, in level as well as in rank, and it is a natural candidate under an economic reading of business time, in the spirit of [Ané and Geman \(2000\)](#). Two qualifications belong with this reading. The elasticity of 1.14 is statistically distinguishable from one ($p = 0.011$), so the match is close rather than exact. And, by Section 6.2, the whole comparison is anchored to one-minute sampling: at a different frequency v would move in a way the model does not reproduce, and the level comparison would have to be repeated there rather than carried across.

Table 10. Does the Fitted Clock Track Observed Quote Activity?

<i>Panel A. Level restriction $v = CV^2(N)$, daily, time-of-day adjusted</i>						
Clock	N	Median CV^2	Median ratio	[p10, p90] of ratio	Ratio in [0.5, 2] (%)	Spearman ρ
NBBO updates	773	0.164	3.12	[0.94, 6.67]	19	0.777
Mid-quote changes	773	0.360	1.24	[0.37, 2.78]	61	0.772
Non-zero seconds	773	0.107	4.74	[1.16, 46.93]	16	0.461
Within-minute RV	773	0.569	0.77	[0.20, 1.52]	68	0.773
<i>Panel B. $\log v = a + b \log CV^2(N) + \text{controls} + \text{fixed effects}$</i>						
Specification	N	b	(t)	$p(b = 1)$	R^2	
<i>Daily, controlling for RV, spread, zero-return share and year effects</i>						
NBBO updates	773	1.594	(21.8)	<0.001	0.511	
Mid-quote changes	773	1.144	(20.2)	0.011	0.491	
Non-zero seconds	773	1.571	(16.7)	<0.001	0.415	
Within-minute RV	773	0.996	(18.1)	0.948	0.477	
<i>Half-hour windows, with day and window-of-day effects</i>						
NBBO updates, no fixed effects	8215	1.697	(40.3)	<0.001	0.444	
NBBO updates, both fixed effects	8215	1.449	(28.9)	<0.001	0.619	
NBBO updates, with controls	8215	0.724	(16.4)	<0.001	0.707	
Mid-quote changes, with controls	8215	0.542	(15.7)	<0.001	0.703	

Notes: Panel A: $CV^2(N) = \text{Var}(N_m)/E(N_m)^2$ for per-minute counts N_m of each candidate clock, after the same leave-one-day-out intraday adjustment applied to returns. If the business clock is market activity, the normalized count is the clock increment, so the model predicts a ratio of one and not merely co-movement. The ratio is $v/CV^2(N)$ by day; the last column is the Spearman correlation of $\log v$ with $\log CV^2$ across days, essentially unchanged when year means are removed. Panel B reports the slope b , its t -statistic against zero in parentheses, the p -value for the model's prediction $b = 1$, and R^2 . Daily rows use the primary estimate (60 s mid-quote returns, time-of-day adjusted, maximum likelihood), control for $\log RV$, the log quoted spread, the zero-return share and year, and have HC3 standard errors; window rows use half-hour windows estimated on 5-second returns, with standard errors clustered by day. Within-minute realized variance is listed for completeness but is not an independent test, because v is identified from the kurtosis of one-minute returns, which variation in within-minute variance itself drives.

6.4 The Event Panel

The case-study files offer three to seven control days per event, so the smallest attainable two-sided randomization p -value in 2008 is 0.125 and no 2008 event can be separated from its controls. The panel provides 166 FOMC statement days, 212 CPI and 209 Employment Situation releases.

The design is the one of Section 4, applied inside the day. Each day is cut into thirteen half-hour windows, and the model is estimated on the 360 five-second mid-quote returns of each window. For every outcome, the abnormal value of a window is its log level minus that day's mean over its own non-announcement windows, minus the mean of the same window over days carrying no such announcement. Both the day's conditions and the intraday pattern are therefore removed, and the reference distribution is the same statistic on days without the announcement. FOMC statements are released at 14:00, so they claim the




Figure 8. Fitted Clock Dispersion against the Dispersion of Quote Counts. *Notes:* One point per day, 773 full-session days. Horizontal axis: $\log CV^2$ of per-minute counts of NBBO updates (left) and of mid-quote changes (right), time-of-day adjusted. Vertical axis: $\log \hat{v}$ from one-minute time-of-day-adjusted mid-quote returns by maximum likelihood. The dashed line is the model’s prediction $v = CV^2(N)$; the solid line is the bivariate fit, whose slope the model puts at one. Table 10 adds controls and fixed effects. Axes are clipped to the 0.5th–99.5th percentiles: the points below the cloud are days on which price discreteness pushes the likelihood towards its lower bound on v .

14:00–14:30 window and, through the press conference, the 14:30–15:00 window. CPI and Employment Situation are released at 08:30, before the open, so their response appears in the first half hour of trading.

One estimation detail governs the sample. The normal inverse Gaussian has no atom at zero. On days when the penny tick is large relative to volatility, half or more of the five-second returns are exactly zero, and the likelihood is then maximized by driving v to its upper bound: the fit converges on essentially every window whose zero-return share is below 0.3 and on none whose share exceeds 0.5. Those degenerate windows are not missing at random—announcement windows are active, so 95% of them converge against 76% of control windows, and the windows lost from the controls are the high- v ones. We therefore report the event study on the 423 days whose median zero-return share over 10:00–13:30 is at most 30%, a property of the day’s tick size relative to its volatility that no announcement touches. Within that sample the fit converges on 99.7% of announcement-capable windows and 99.8% of the rest. Table 11 gives both samples.

Three patterns emerge, and they bear on the two objects the framework is built to distinguish.

First, clock dispersion responds to the FOMC statement, and we find little evidence of a response to the other announcements. In the 14:00–14:30 window v is 0.78 log points above its counterfactual, a factor of about 2.2, with a randomization p -value at the floor of the test. In the press-conference window that follows it is 0.09 log points above, which is not distinguishable from zero ($p = 0.50$). On the full 779-day sample the corresponding press-conference estimate is negative (-0.57); because the likelihood degenerates on high-discreteness days in a way that differs between announcement and control windows, we

place more weight on the 423-day sample for v , and neither sample shows a positive press-conference response comparable to that in the statement window. On CPI days the first half hour has v 0.34 log points *below* its counterfactual ($p = 0.018$), and on Employment Situation days it is unmoved ($p = 0.84$). This is more precise than the case-study pattern of Section 5, in which the excess appeared concentrated in the hour containing the statement and the press conference: with 166 statement days the response appears to be located in the half hour of the statement itself.

Second, volatility and quoted spreads behave differently. Realized volatility is 1.62 and 1.66 log points above its counterfactual in the two FOMC windows—a factor of about five in both—and quoted spreads are 0.28 and 0.26 log points above, about 30% wider, while displayed depth falls by about a quarter. Both are also elevated, modestly but detectably, on CPI and Employment Situation days. The magnitude of uncertainty therefore appears to rise around macroeconomic news and to remain high through the press conference, whereas the irregularity of its timing rises only at the discrete moment of the policy statement. This pattern is consistent with the distinction the framework is designed to draw.

Third, the dispersion of observed quote activity does not move detectably with either. The coefficient of variation of five-second NBBO update counts within the window is not distinguishable from its counterfactual in the FOMC windows ($p = 0.11$ and $p = 0.85$). Whatever raises v at 14:00, it does not appear to be a change in how unevenly quote updates arrive; Section 6.3 takes up the observable counterpart of the clock. Figure 9 shows the three patterns together across the trading day: v rises in the statement window and is back to the control level in the next one, while the quoted spread rises at 14:00 and stays elevated to the close.

The same comparison at the level of the whole day illustrates the paper’s main distinction, and it extends the case-study pattern of Section 5 to 166 statement days. Estimated on the day’s 390 one-minute returns, v on an FOMC statement day is about 4.1 times its median on days without a statement ($p = 0.0001$), while daily realized volatility is 1.17 times its median ($p = 0.0005$) and the daily quoted spread and depth show no detectable difference ($p = 0.23$ and 0.32 ; Table 12). Two points deserve care. The result does not appear to be a mechanical consequence of the window result: the median of v across the day’s thirteen half-hour windows is close to unchanged (1.15 times, $p = 0.68$), which is what one would expect if a single window in thirteen were anomalous. And the comparison group is not a neutral baseline—the announcement-free days are drawn from the neighborhoods of the 2008 and 2025 case studies and are therefore unusually volatile, which is consistent with CPI and employment days showing realized volatility *below* it. The more consistent comparison is the contrast within the same days against the same group: a factor of about four in timing irregularity against a factor of 1.17 in magnitude. A model-free counterpart points the same way. The median sample excess kurtosis of the day’s time-of-day-adjusted one-minute mid-quote returns is 6.7 on statement days against 1.1 on other days, and 10.2 against 2.3 at 15-second sampling (permutation $p = 0.0001$ in both cases; 92% of statement days lie above the median of the other days). This does not rely on the likelihood fit or on the degeneracy filter, although, because v is identified from kurtosis, it is a consistency check rather than an independent test.

Consistent with the window results, the day-level effect is associated with policy statements and not with macroeconomic releases. On CPI days daily v is not distinguishable from its comparison level (0.96, $p = 0.36$); on employment days it is lower (0.78,

$p = 0.0008$), as is realized volatility. The clock *share* of variance shows no detectable difference on any of the three ($p \geq 0.39$), in line with Section 6.1, where timing irregularity appears to reside in the tails of the return distribution rather than in the variance.

Finally, the measured size of the news does not appear to explain the size of the response. We regress the same abnormal values on the Jarociński and Karadi (2020) monetary-policy surprise—the high-frequency identification approach of Kuttner (2001) and Gürkaynak et al. (2005), with the reassessment of Bauer and Swanson (2023)—on its monetary-policy and central-bank-information components, and on the S&P 500 surprise. In the 423-day sample no coefficient has $|t|$ above 1.4 and no R^2 exceeds 0.08, for v , realized volatility, the quoted spread or update dispersion alike. On the full 779-day sample the absolute monetary-policy component is associated with realized volatility ($t = 2.6$, $R^2 = 0.07$), whereas no comparable association appears for v (largest $|t|$ of 1.3). With 71 statement days in the low-discreteness sample these estimates carry limited power; they are consistent with the occurrence of a statement moving these series, and offer little evidence on whether its measured surprise explains how much.




Figure 9. Abnormal Clock Dispersion and Quoted Spread by Half-Hour Window. *Notes:* Median abnormal value—log level net of the day’s own non-announcement level and of the window’s time-of-day level—by window, on FOMC statement days against all other days, on the 423 low-discreteness days described in the text. The dotted line marks 14:00, when the statement is released. Left: clock dispersion v from five-second returns; it returns to the control level in the next window. Right: the quoted spread, which stays elevated through the close.

6.5 Does Clock Dispersion Predict Subsequent Market Conditions?

A financial-stability reading of this framework would be a forecasting claim: that timing irregularity today signals deterioration in liquidity tomorrow. The case studies are too short to examine it, and the panel allows a test. We estimate

$$Y_{t+h} = \alpha_h + \beta_h Z_t + \gamma'_h X_t + \varepsilon_{t+h}, \quad (16)$$

Table 11. The Event Panel: Abnormal Half-Hour Windows

Event	Window	Outcome	N	Controls	Effect	z	p	All days
FOMC statement	14:00–14:30	Clock dispersion v	72	344	0.784	1.27	<0.001	0.379
		Realized volatility	72	346	1.616	3.77	<0.001	1.950
		Quoted spread	72	346	0.275	6.10	<0.001	0.153
		Depth at the NBBO	72	346	−0.300	−2.31	<0.001	−0.374
		Dispersion of updates	72	346	0.072	0.55	0.109	−0.111
	14:30–15:00	Clock dispersion v	72	344	0.094	0.01	0.504	−0.567
		Realized volatility	72	346	1.660	3.54	<0.001	1.577
		Quoted spread	72	346	0.260	3.49	<0.001	0.116
		Depth at the NBBO	72	346	−0.322	−2.52	<0.001	−0.303
		Dispersion of updates	72	346	−0.007	−0.15	0.851	−0.313
CPI release	09:30–10:00	Clock dispersion v	97	326	−0.336	−0.31	0.018	−0.277
		Realized volatility	97	326	0.158	0.55	0.013	0.176
		Quoted spread	97	326	0.022	0.28	0.116	0.007
		Depth at the NBBO	97	326	−0.020	−0.17	0.396	−0.012
		Dispersion of updates	97	326	−0.131	−0.33	0.003	−0.104
Employment Situation	09:30–10:00	Clock dispersion v	104	319	−0.028	0.10	0.843	−0.260
		Realized volatility	104	319	0.121	0.30	0.062	0.279
		Quoted spread	104	319	0.033	0.36	0.016	0.014
		Depth at the NBBO	104	319	−0.040	−0.30	0.084	−0.056
		Dispersion of updates	104	319	0.017	0.05	0.695	−0.100

Notes: Each day is cut into thirteen half-hour windows. The abnormal value of a window is its log outcome minus that day's mean over its own non-announcement windows minus the mean of the same window over days without that announcement, which removes both the day's conditions and the intraday pattern. Effect is the difference in means between announcement days and days without the announcement, in log points; z is the robust standardized distance of the announcement-day value from the control median, and p the two-sided randomization p -value with 9,999 re-allocations. CPI and Employment Situation are released at 08:30, before the open, so their window is the first half hour of trading. The sample is the 423 days whose median zero-return share over 10:00–13:30 is at most 30%, a property of the day's tick size relative to its volatility that no announcement affects. The restriction matters only for v : the normal inverse Gaussian has no atom at zero, so on days when discreteness dominates the 5-second returns the likelihood runs to the upper bound on v , and it does so more often in control windows than in announcement windows. Within the sample the fit converges in 99.7% and 99.8% of the two groups. The last column gives the same difference on all 779 days.

Table 12. The Event Panel: Whole-Day Comparison

Event	Outcome	N	Controls	Ratio	z	p
FOMC statement	Clock dispersion v	166	607	4.12	1.73	<0.001
	Median v across windows	165	605	1.15	0.15	0.676
	Realized volatility	166	607	1.17	0.27	<0.001
	Quoted spread	166	607	1.03	0.04	0.227
	Depth at the NBBO	166	607	0.86	-0.09	0.318
	Clock share of variance	166	607	0.79	-0.12	0.553
CPI release	Clock dispersion v	212	561	0.96	-0.04	0.359
	Median v across windows	211	559	1.01	0.01	0.025
	Realized volatility	212	561	0.79	-0.43	<0.001
	Quoted spread	212	561	0.95	-0.08	0.745
	Depth at the NBBO	212	561	0.84	-0.12	0.194
	Clock share of variance	212	561	1.05	0.03	0.169
Employment Situation	Clock dispersion v	206	567	0.78	-0.22	<0.001
	Median v across windows	205	565	1.16	0.16	0.804
	Realized volatility	206	567	0.87	-0.26	0.002
	Quoted spread	206	567	0.95	-0.08	0.711
	Depth at the NBBO	206	567	0.76	-0.18	0.141
	Clock share of variance	206	567	1.24	0.11	0.393

Notes: Each entry compares the whole day on announcement days with days without that announcement, using one-minute returns on all 773 full-session days and no discreteness restriction (at one-minute sampling the zero-return share has a median of 5.6% and every fit converges). Ratio is the announcement-day median divided by the median on days without the announcement, so 1.00 indicates no difference; z and p are defined as in Table 11. The announcement-free days are drawn from the neighborhoods of the case studies and are unusually volatile, which is consistent with realized volatility on CPI and employment days falling below them; the comparison of interest is therefore between rows within an event rather than between an event and one. The median of v across a day's thirteen windows barely moves on statement days while the whole-day estimate is four times as large, which is consistent with a burst confined to one window in thirteen.

with $Z_t \in \{\log v_t, s_t^C, \log \text{SUI}_t\}$, controls X_t containing the current log realized volatility and log quoted spread, horizons $h \in \{1, 5, 10, 22\}$ trading days, and outcomes covering quoted-spread widening, realized volatility, the change in the VIX, Amihud illiquidity (Amihud, 2002), the closing spread, absolute returns, the OFR Financial Stress Index and the two-year Treasury yield. Standard errors are Newey and West (1987) with $2h$ lags, the truncation stated in the table; Hodrick (1992) is a natural alternative for overlapping horizons, and we rely instead on the bootstrap described next, which addresses the same problem without a lag choice. Because both the predictor and the outcomes are persistent and the horizons overlap, the conventional t -statistic may not be reliable here (Stambaugh, 1999; Valkanov, 2003); we therefore report a stationary-bootstrap p -value under the null that Z adds nothing beyond X , constructed by resampling the residuals of Y on X in blocks while leaving Z in its observed order. Out-of-sample forecasts are expanding-window and rolling-origin, scored against the $\{\text{RV}, \text{spread}\}$ baseline by Clark and West (2007) and Diebold and Mariano (1995). Table 13 reports the results for $\log v$.

We find little evidence of predictive content. Across 96 in-sample specifications, 14 reach $p < 0.05$ where about five would be expected by chance, and the incremental R^2 over the baseline has a median of 0.001 and a maximum of 0.039. Out of sample, 9 of 32 forecasts improve on the baseline at all, five have a Clark–West statistic above 1.96, and none has a Diebold–Mariano statistic above 1.96. We do not find a horizon or an outcome at which clock dispersion delivers a usable forecast of liquidity deterioration.

Where coefficients are significant, they mostly do not have the sign that an early-warning interpretation would need. The most consistent result in the table is that $\log v$ is associated with *lower* future realized volatility: -0.046 at one day ($t = -4.70$, bootstrap $p = 0.001$), -0.047 at ten days and -0.036 at twenty-two, with the same negative sign for absolute returns. Given that the regression already controls for today’s realized volatility, and that v and realized volatility are positively related contemporaneously, this is the pattern that mean reversion would produce rather than the pattern that a stress signal would produce. The one outcome with the sign an early-warning story would want is the VIX, which rises after high- v days at ten and twenty-two days ($\beta = 0.010$ and 0.011 , bootstrap $p = 0.036$ and 0.052); the incremental R^2 is 0.006 and 0.004, and out of sample both forecasts are worse than the baseline.

Three features of the sample bound what this can settle. The panel is an event sample, so these regressions condition on announcement days and case-study neighborhoods; restricting to the 200 days that carry no announcement, and adding announcement-type fixed effects, does not materially change the results. The CRSP daily file we use ends on 31 December 2024, so the Amihud and closing-spread rows lose the 2025 days. And the NBBO outcomes exist only on pulled days, so their horizon is the first pulled day at least h days ahead, which we require to be within five trading days of the target. With 773 days, horizons out to a month and eight outcomes, a strong predictive relation between clock dispersion and liquidity stress would likely have appeared in this table, although modest effects cannot be ruled out.

Table 13. Clock Dispersion and Subsequent Market Conditions

Outcome	h	N	In-sample					Out-of-sample	
			β	(t_{NW})	p	Boot. p	ΔR^2	R^2_{OOS}	CW
NBBO quoted spread	1	438	-0.0032	(-0.38)	0.704	-	0.0008	-0.0122	-1.47
	5	499	-0.0003	(-0.04)	0.967	-	0.0000	-0.0110	-1.60
	10	462	-0.0051	(-0.54)	0.588	-	0.0017	-0.0095	-0.72
	22	532	-0.0054	(-0.82)	0.412	-	0.0012	-0.0071	-0.94
Realized volatility	1	438	-0.0455	(-4.70)	<0.001	0.001	0.0386	0.0175	2.85
	5	499	-0.0338	(-2.09)	0.037	0.059	0.0121	0.0153	1.57
	10	462	-0.0469	(-3.25)	0.001	0.007	0.0179	0.0271	1.89
	22	532	-0.0357	(-3.24)	0.001	0.013	0.0072	0.0096	2.35
VIX (log change)	1	773	0.0009	(0.32)	0.748	-	0.0002	-0.0058	0.27
	5	773	0.0050	(0.96)	0.337	-	0.0024	-0.0197	0.34
	10	773	0.0098	(2.30)	0.022	0.036	0.0056	-0.0066	1.19
	22	773	0.0111	(2.12)	0.034	0.052	0.0036	-0.0028	1.50
Amihud illiquidity	1	656	-0.0190	(-0.67)	0.506	-	0.0006	-0.0027	-0.78
	5	655	-0.0325	(-1.05)	0.294	-	0.0018	-0.0019	0.32
	10	650	-0.0205	(-0.64)	0.523	-	0.0007	-0.0037	-0.78
	22	640	-0.0519	(-2.39)	0.017	0.037	0.0044	0.0028	1.14
Closing spread	1	584	-0.0097	(-0.94)	0.347	-	0.0008	-0.0049	-0.84
	5	587	-0.0114	(-1.14)	0.256	-	0.0010	-0.0095	0.15
	10	580	0.0015	(0.29)	0.772	-	0.0000	-0.0055	-0.31
	22	565	0.0013	(0.18)	0.854	-	0.0000	-0.0076	-1.68
Absolute return	1	659	-0.0011	(-3.19)	0.001	0.001	0.0131	0.0249	2.88
	5	656	-0.0012	(-2.30)	0.021	0.018	0.0175	0.0459	2.90
	10	652	-0.0002	(-0.55)	0.580	-	0.0008	-0.0154	-1.65
	22	641	-0.0005	(-2.30)	0.021	0.031	0.0026	0.0075	2.88
OFR Financial Stress Index	1	772	-0.0187	(-1.37)	0.172	-	0.0032	0.0006	0.66
	5	773	0.0553	(1.42)	0.155	-	0.0053	-0.0804	0.15
	10	771	0.0964	(1.46)	0.143	-	0.0075	-0.0793	0.80
	22	773	0.1068	(1.14)	0.253	-	0.0030	-0.0124	1.34
2-year Treasury yield	1	750	-0.0033	(-1.17)	0.244	-	0.0051	-0.0002	0.66
	5	760	-0.0050	(-1.61)	0.108	0.116	0.0035	-0.0100	0.58
	10	764	-0.0007	(-0.19)	0.851	-	0.0000	-0.0075	-1.51
	22	757	0.0004	(0.05)	0.957	-	0.0000	-0.0123	-1.98

Notes: $Y_{t+h} = \alpha + \beta \log v_t + \gamma' X_t + e$ with $X_t = \{\log RV_t, \log \text{spread}_t\}$ and h in trading days. $\log v$ is the primary daily estimate (60 s mid-quote returns, time-of-day adjusted, maximum likelihood). Newey–West t -statistics with $2h$ lags are in parentheses. Boot. p is a stationary-bootstrap p -value under no predictability, computed where $|t_{NW}| > 1.5$; it is the relevant reference given the persistence of both sides (Stambaugh, 1999; Valkanov, 2003). ΔR^2 is the gain over the $\{RV, \text{spread}\}$ baseline; R^2_{OOS} and the Clark–West statistic (CW) come from expanding-window rolling-origin forecasts against that baseline (Clark and West, 2007). Only $\log v$ is reported; regressions with the clock share and with $\log \text{SUI}$ as predictors are summarized in the text, and the SUI level equals the sample volatility under the exactly identified moment estimator. Amihud illiquidity and the closing spread come from CRSP, whose daily file used here ends on 31 December 2024, so those rows lose the 2025 days.

7 Interpretation, Limitations, and Future Work

The central empirical question is whether the timing of quote revisions provides a distinct description of short-lived liquidity stress. The evidence is consistent with a limited version of that proposition. The fitted clock is related to the dispersion of mid-quote revisions, and its largest within-day response occurs around FOMC statements, where the change is concentrated in the statement window. The result is descriptive rather than causal and does not imply that the clock provides a general measure of financial instability.

7.1 What the Evidence Suggests

The first supporting result concerns measurement. Under the exactly identified moment estimator, the SUI is the sample standard deviation by construction (Remark 1), and the identity holds to machine precision across the panel. The likelihood estimator produces a very similar result. Thus, the paper does not propose a new volatility measure. The information of interest lies in the decomposition and, in particular, in the clock-dispersion component v .

The second supporting result concerns the observable counterpart of v . Across the panel, v is positively related to the dispersion of mid-quote revision counts, with rank correlations near 0.77 after accounting for year, day, and time-of-day effects and controlling for volatility and the quoted spread. The relation is weaker for raw quote traffic, which is smoother than the fitted clock. This evidence is consistent with interpreting v as a measure of the irregularity of price-relevant revisions rather than of quoting activity in general.

The main empirical result concerns liquidity stress within event windows. On FOMC statement days, clock dispersion rises sharply around the statement and then returns toward its usual level, while realized volatility and quoted spreads remain elevated more broadly around the announcement. The case-study evidence and the longer panel point in the same direction. The timing measure therefore provides a more localized description of the episode than daily volatility alone. The comparison is not a causal event study: policy statements occur at a fixed time, so the response to the event cannot be completely separated from the response to that time of day.

7.2 Two Qualifications

The first qualification concerns the model. A Lévy subordinator implies that excess kurtosis falls as $1/\tau$ with the sampling interval. In the data, the decline is closer to $\tau^{-0.46}$ over the range examined in Section 6.2. This discrepancy suggests that the independent-increment clock is an imperfect description of the data. The identification result remains a statement about the information contained in return cumulants, but the economic interpretation of v as the variance of a Lévy clock should be treated cautiously. We therefore report the sampling interval with the estimates and do not treat v as a frequency-invariant parameter.

The second qualification concerns forecasting. Across the liquidity outcomes and horizons examined, we find little evidence that clock dispersion predicts subsequent market conditions once conventional measures such as realized volatility and the quoted spread are included. The evidence therefore does not support treating v as an early-warning indicator. This is consistent with the paper’s narrower objective: describing the timing of

liquidity stress within an observed event rather than predicting future stress.

7.3 Implications for Financial Stability

The financial-stability relevance of the result is primarily one of measurement. A daily volatility measure aggregates price revisions over the day and can therefore obscure whether the adjustment was spread broadly across the session or concentrated in a short interval. The clock-dispersion measure provides a way to distinguish these patterns. Around FOMC statements, the data show a pronounced increase in the irregularity of price revisions in the statement window, while the corresponding increase in conventional market measures extends beyond that window.

This distinction should not be interpreted as evidence of a new supervisory indicator. The panel does not show that clock dispersion contains incremental information about future liquidity conditions, and the event design cannot establish causality. Rather, the result suggests that high-frequency monitoring of quote-revision timing can complement measures based on volatility and spreads when the objective is to describe how a liquidity-stress episode unfolds within the trading day.

7.4 Limits of the Evidence

Several features limit the scope of the result. The panel is concentrated on event days and neighborhoods of stress episodes, so its unconditional moments should not be interpreted as representative of ordinary trading days. SPY also has a one-cent tick and typically a one-tick spread, which limits identification of bid-ask asymmetry at the sampling frequency used for the model. The normal inverse Gaussian distribution has no atom at zero, creating estimation problems when discrete price changes dominate at very high frequencies. Finally, the TAQ sample does not yet provide a complete 2026 counterpart for all of the index-quote case studies.

The event comparisons also have an important design limitation. FOMC, CPI, and employment releases occur at scheduled times, so the response of the clock cannot be interpreted as a causal event effect. The use of same-clock-time controls reduces the role of predictable intraday variation but does not remove this identification problem.

7.5 Directions for Future Work

The most direct extension is to develop a clock with dependent increments. The aggregation evidence suggests that price-revision activity persists across sampling intervals in a way that a Lévy subordinator does not capture. A model with persistent or regime-dependent activity could provide a more stable interpretation of clock dispersion across frequencies and could allow the measurement framework to be evaluated on instruments with richer liquidity variation.

A second extension would use markets with more variation in spreads and depth, together with sub-second quote data. Such data would permit a more direct examination of whether timing irregularity differs across the two sides of the market and how it relates to effective spreads and order-book depth.

A third extension is to incorporate the discrete nature of quoted prices explicitly. Because a large fraction of the highest-frequency observations are zero returns, a rounded or censored normal-variance mixture could allow the framework to use shorter sampling intervals without treating price discreteness as a nuisance. These extensions would test whether the central empirical relation between revision timing and liquidity stress survives a richer description of high-frequency market activity.

8 Conclusion

This paper asks whether the timing of quote revisions provides information about short-lived episodes of market liquidity stress that is not captured by the magnitude of price movements. We develop a time-changed return model that separates conventional return variation from dispersion in a stochastic business-time clock and apply it to high-frequency S&P 500 and SPY quote data.

The central empirical result is that clock dispersion is related to the dispersion of mid-quote revisions and rises sharply around FOMC statements. The increase is concentrated in the statement window, providing a more localized description of the episode than daily volatility alone. Similar responses are not observed around CPI and Employment Situation releases. The comparisons are descriptive rather than causal because the announcements occur at fixed times of day.

The remaining results mainly establish the scope of this finding. The model is identified only in its single-subordinator form from fourth-order return cumulants, and its uncertainty index coincides with sample volatility under the moment estimator. The data also reject the aggregation behavior implied by an independent-increment Lévy clock, and we find little evidence that clock dispersion predicts subsequent liquidity conditions. These qualifications argue against interpreting the measure as a new volatility index or an early-warning indicator.

The evidence instead supports a narrower interpretation. During a liquidity-stress episode, the timing and concentration of quote revisions can contain information that is obscured when market adjustment is summarized only by aggregate volatility. Whether this timing dimension can be developed into a broader measure of market stability requires a model that better accommodates persistent activity, price discreteness, and richer variation in spreads and depth. We leave those extensions for future work.

Disclosure Statement

The authors report there are no competing interests to declare.

Data Availability Statement

SPX best bid and ask quotes are obtained from Bloomberg’s $\langle GP \rangle$ function at the minute-level. The SPY National Best Bid and Offer is from the WRDS TAQ millisecond database and the daily stock file from CRSP, both subject to subscription; the code that draws them, including the server-side aggregation into one-second bars, is part of

the replication package, so a subscriber can reproduce the panel from source. The event calendars (Federal Reserve press releases, Bureau of Labor Statistics release archives), the high-frequency monetary-policy surprises of [Jarociński and Karadi \(2020\)](#), the outcome series (CBOE VIX, OFR Financial Stress Index, Treasury par yield curve) and the replication code are publicly available.

Definitions

Returns and volatility. For day d with one-minute log returns $r_{d,m}$, $m = 1, \dots, 389$, realized volatility is $\text{RV}_d = \sqrt{390 \times 252} \left(\frac{1}{389} \sum_m r_{d,m}^2 \right)^{1/2}$ and historical volatility HV_d is $\sqrt{390 \times 252}$ times the sample standard deviation of $r_{d,m}$. The SUI is annualized by the same factor.

Intraday profile. For each period and side, $f(m)$ is proportional to the square of the average across control days of $|r_{d,m}|/|r_{d,\cdot}|$, smoothed with a centered 11-minute moving average after the first five minutes and normalized to mean one; time-of-day-adjusted returns are $r_{d,m}/\sqrt{f(m)}$, with each control day excluded from its own profile.

Randomization inference. For an event value e and control values $c_1, \dots, c_n$ with median $\bar{c}$, the two-sided p -value is $(1 + \#\{k : |c_k - \bar{c}| \geq |e - \bar{c}|\})/(1 + n)$.

References

- Aït-Sahalia, Y. and Jacod, J. (2010). Is brownian motion necessary to model high-frequency data? *Annals of Statistics*, 38(5):3093–3128.
- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.
- Andersen, T. G. and Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2–3):115–158.
- Ané, T. and Geman, H. (2000). Order flow, transaction clock, and normality of asset returns. *Journal of Finance*, 55(5):2259–2284.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.
- Bali, T. G., Hu, J., and Murray, S. (2021). Option implied volatility, skewness, and kurtosis and the cross-section of expected stock returns. *Journal of Accounting and Economics*, 74(1).
- Ballotta, L., Eberlein, E., Schmidt, T., and Zeineddine, R. (2024). Multivariate additive subordination with applications in finance. *European Journal of Operational Research*, 317(2):469–484.
- Banerjee, S. and Kremer, I. (2010). Disagreement and learning: Dynamic patterns of trade. *Journal of Finance*, 65(4):1269–1302.

- Barndorff-Nielsen, O. E. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics*, 24(1):1–13.
- Barndorff-Nielsen, O. E. (1998). Processes of normal inverse Gaussian type. *Finance and Stochastics*, 2(1):41–68.
- Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155.
- Bekaert, G., Hoerova, M., and Lo Duca, M. (2013). Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788.
- Biais, B., Glosten, L., and Spatt, C. (2005). Market microstructure: A survey of micro-foundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2):217–264.
- Bianchi, M. L., Rachev, S. T., Kim, Y. S., and Fabozzi, F. J. (2011). Tempered infinitely divisible distributions and processes. *Theory of Probability and Its Applications*, 55(1):2–26.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.
- Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11):4463–4492.
- Bonart, J. and Gould, M. D. (2018). Latency and liquidity provision in a limit order book. *Quantitative Finance*, 18(3):453–469.
- Bouchaud, J.-P., Mézard, M., and Potters, M. (2002). Statistical properties of stock order books: Empirical results and models. *Quantitative Finance*, 2(4):251–256.
- Brandt, M. W. and Kavajecz, K. A. (2004). Price discovery in the U.S. treasury market: The impact of orderflow and liquidity on the yield curve. *Journal of Finance*, 59(6):2623–2654.
- Briola, A., Vidal-Tomás, D., and Aste, T. (2024). Deep limit order book forecasting: A microstructural guide. *Quantitative Finance*, 25(6):1101–1131.
- Buraschi, A., Trojani, F., and Vedolin, A. (2014). When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia. *Journal of Finance*, 69(1):101–137.
- Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.
- Carlin, B. I., Longstaff, F. A., and Matoba, K. (2014). Disagreement and asset prices. *Journal of Financial Economics*, 114(2):226–238.
- Carr, P., Chang, E., and Madan, D. B. (1998). The variance gamma process and option pricing. *European Finance Review*, 2(1):79–105.
- Carr, P., Geman, H., Madan, D. B., and Yor, M. (2002). The fine structure of asset returns: An empirical investigation. *Journal of Business*, 75(2):305–332.

- Carr, P. and Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341.
- Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41(1):135–155.
- Clark, T. E. and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1):291–311.
- Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236.
- Cont, R., Stoikov, S., and Talreja, R. (2010). A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563.
- Cont, R. and Tankov, P. (2004). *Financial Modelling with Jump Processes*. Chapman and Hall/CRC, Boca Raton, FL.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.
- Drechsler, I. and Yaron, A. (2011). What’s vol got to do with it. *Review of Financial Studies*, 24(1):1–45.
- Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.
- Easley, D., Lopez de Prado, M. M., and O’Hara, M. (2011). The microstructure of the “flash crash”: Flow toxicity, liquidity crashes, and the probability of informed trading. *Journal of Portfolio Management*, 37(2):118–128.
- Engle, R. F. (2000). The econometrics of ultra-high-frequency data. *Econometrica*, 68(1):1–22.
- Engle, R. F. and Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66(5):1127–1162.
- Geman, H., Madan, D. B., and Yor, M. (2001). Time changes for Lévy processes. *Mathematical Finance*, 11(1):79–96.
- Glosten, L. R. and Harris, L. E. (1988). Estimating the components of the bid/ask spread. *Journal of Financial Economics*, 21(1):123–142.
- Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100.
- Gould, M. D., Porter, M. A., Williams, S., McDonald, M., Fenn, D. J., and Howison, S. D. (2013). Limit order books. *Quantitative Finance*, 13(11):1709–1742.

- Gürkaynak, R. S., Sack, B., and Swanson, E. T. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1):55–93.
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, 46(1):179–207.
- Hasbrouck, J. (2010). The best bid and offer: A short note on programs and practices. *Journal of Financial Markets*, 13(3):289–297.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343.
- Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies*, 5(3):357–386.
- Huang, J.-Z. and Wu, L. (2004). Specification analysis of option pricing models based on time-changed Lévy processes. *Journal of Finance*, 59(3):1405–1440.
- Huang, R. D. and Stoll, H. R. (1997). The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10(4):995–1034.
- Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.
- Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.
- Kelly, B. T. and Jiang, H. (2014). Tail risk and asset prices. *Review of Financial Studies*, 27(10):2841–2871.
- Kim, Y. S., Rachev, S. T., Bianchi, M. L., and Fabozzi, F. J. (2010). Tempered stable and tempered infinitely divisible GARCH models. *Journal of Banking and Finance*, 34(9):2096–2109.
- Korajczyk, R. A. and Murphy, D. (2019). High-frequency market making to large institutional trades. *Review of Financial Studies*, 32(3):1034–1067.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.
- Li, J. and Zinna, G. (2018). The variance risk premium: Components, term structures, and zeta-risk. *Journal of Business and Economic Statistics*, 36(2):199–216.
- Ludvigson, S. C., Ma, S., and Ng, S. (2021). Uncertainty and business cycles: Exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics*, 13(4):369–410.

- Madan, D. B. and Seneta, E. (1990). The variance gamma model for share market returns. *Journal of Business*, 63(4):511–524.
- Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3(3):205–258.
- Mandelbrot, B. and Taylor, H. M. (1967a). On the distribution of stock price differences. *Operations Research*, 15(6):1057–1062.
- Mandelbrot, B. B. and Taylor, H. M. (1967b). On the distribution of stock price differences. *Operations Research*, 15(6):1057–1062.
- Moinas, S., Nguyen, M., and Vuillemeys, G. (2020). Equilibrium in the limit order book: Does tick size matter? *Journal of Finance*, 75(2):603–640.
- Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.
- Ntakaris, A., Magris, M., Kannianen, J., Gabbouj, M., and Iosifidis, A. (2018). Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods. *Journal of Forecasting*, 37(8):852–866.
- O’Hara, M. (1995). *Market Microstructure Theory*. Blackwell, Cambridge, MA.
- Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, 63(1):3–50.
- Rachev, S. T., Kim, Y. S., Bianchi, M. L., and Fabozzi, F. J. (2011). *Financial Models with Lévy Processes and Volatility Clustering*. John Wiley & Sons, Hoboken, NJ.
- Rachev, S. T. and Mitnik, S. (2000). *Stable Paretian Models in Finance*. John Wiley & Sons, Chichester.
- Riordan, R. and Storkenmaier, A. (2012). Latency, liquidity and price discovery. *Journal of Financial Markets*, 15(4):416–437.
- Semeraro, P. (2022). Multivariate tempered stable additive subordination for financial models. *Mathematics*, 10(3):363.
- Shirvani, A., Mitnik, S., Lindquist, W. B., and Rachev, S. T. (2024). Bitcoin volatility and intrinsic time using double-subordinated Lévy processes. *Risks*, 12(5):82.
- Shirvani, A., Rachev, S. T., and Fabozzi, F. J. (2020). Multiple subordinated modeling of asset returns: Implications for option pricing. *Econometric Reviews*, 40(3):290–319.
- Shirvani, A., Stoyanov, S., Fabozzi, F. J., and Rachev, S. T. (2021). Equity premium puzzle or faulty economic modelling? *Review of Quantitative Finance and Accounting*, 56(4):1329–1342.
- Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics*, 54(3):375–421.

- Valkanov, R. (2003). Long-horizon regressions: Theoretical results and applications. *Journal of Financial Economics*, 68(2):201–232.
- Whaley, R. E. (2000). The investor fear gauge. *Journal of Portfolio Management*, 26(3):12–17.
- Zhang, Z., Zohren, S., and Roberts, S. (2019). DeepLOB: Deep convolutional neural networks for limit order books. *IEEE Transactions on Signal Processing*, 67(11):3001–3012.