

Pre-FOMC Uncertainty Accumulation: Evidence from 0DTE Options

Yuhan Song^{*†}

September 15, 2026

ABSTRACT

Using 0DTE SPX options, I show that the option-implied uncertainty of the prospective shock accumulates rather than resolves before scheduled FOMC announcements. The implied uncertainty monotonically rises from 10:00 ET to the 14:00 ET release and largely collapses at the announcement. Both the level and the accumulation of the implied uncertainty strongly forecast realized volatility in any window after the announcement time, but not prior to it. The implied skewness is uniformly negative, varies little, and does not predict returns. Immediately after the introduction of Wednesday-expiring 0DTE options, the pre-announcement drift documented in the literature disappears, indicating that the drift reflected compensation for previously unobservable event uncertainty. The same accumulation-and-release pattern appears around FOMC Minutes and the closing auction, scaled to their information content. The 0DTE option market prices any regularly anticipated concentration of price discovery in the same accumulating fashion, whether macroeconomic or microstructural.

JEL Classification: G12; G13; G14.

Keywords: FOMC; Event Uncertainty; 0DTE; Closing Auction.

^{*}HEC Montréal, yuhan.song@hec.ca.

[†]I'm deeply grateful to my supervisors, Christian Dorion and Piotr Orłowski, for their invaluable guidance and generous support throughout this project.

1. Introduction

The Federal Open Market Committee (FOMC) announcement is a scheduled event whose timing is known in advance, while its policy content is not. For two decades, U.S. equity markets earned most of their daily expected return in the window before each release, at unusually low realized volatility (Lucca and Moench, 2015; Hu et al., 2022). The literature has not converged on a unified explanation. Cieslak, Morse, and Vissing-Jorgensen (2019) attribute it to leakage of the policy stance through informal Fed communication between meetings—i.e., early revelation of the decision itself. Johannes, Kaeck, and Seeger (2026), however, directly measure option-implied FOMC event risk and find no systematic pre-announcement trend. Event uncertainty does not systematically resolve before FOMC announcements, but neither does it systematically build. I show that it does. The uncertainty of the event impact *accumulates* monotonically until the announcement, and resolves only *after* the information is delivered. At the same time, the pre-announcement return drift that made this window famous is gone in recent years. The break dates to a single announcement—the first one after a same-day option expiration was introduced.

To measure this event uncertainty, I use intraday zero-days-to-expiration (0DTE) S&P 500 option prices to extract the option-implied uncertainty and directional tilt of the prospective market reaction: the expected event volatility and the expected event skewness. The expected event volatility captures the consensus size of the upcoming stock market move, and a larger expected size means the market sees a wider distribution of possible realizations. The rising implied volatility before the event therefore indicates rising uncertainty about its impact. These option-implied expectations are one half of the design; the other is the realization. I measure the uncertainty realization using high-frequency realized volatility of index futures, whose discontinuous component over the announcement window captures the one-shot repricing through which the market absorbs the policy content. Realized volatility therefore serves as the outcome variable, providing the ex post counterpart to the event uncertainty implied by option prices. First, I show the expected event uncertainty follows a clear accumulation-and-release path around FOMC announcements. Through the morning it rises monotonically by 30%, from a return volatility of 0.60% at 10:00 ET to 0.78% at the release, the expected announcement impact on the index level that realizes within minutes. It then collapses by more than half, rebuilds ahead of the press conference, and collapses again

once the press conference begins. In variance units, 71% of the priced event variance resolves at the statement and the remaining 29% at the press conference. It matches the realizations, in which realized discontinuous volatility spikes nearly five times higher at the statement than at the start of the press conference. The minority press-conference share predicts a pattern [Johannes, Kaeck, and Seeger \(2026\)](#) flag as curious: buying straddles loses significantly over the statement yet earns insignificant returns over the press conference. The premium is paid where the uncertainty resolves, a puzzle under their majority press-conference share but a prediction under mine. Predictive regressions confirm that this expected event uncertainty forecasts realized volatility only when the target window contains the announcement itself—not before, with R^2 of 78%–85% at the windows covering the statement and the press conference. The buildup is not a rising risk premium added to forecast content formed early in the day. As uncertainty accumulates over the pre-announcement hours, the later increments have predictive power for the event-window realized volatility, controlling for initial uncertainty levels. These results indicate that event uncertainty does not resolve in advance; instead the option market prices the prospective shock, which is realized only at the announcement. Moreover, the expected event skewness holds a stable negative tilt through the pre-announcement hours and has no predictive power for market returns at almost every horizon. The persistent negative tilt is consistent with the pricing of downside insurance ([Beckmeyer, Branger, and Grünthaler, 2021](#)), rather than a directional forecast of the announcement outcome. Thus, as event uncertainty builds, it does so primarily through the scale of the prospective shock rather than through increasingly negative asymmetry.

Second, I document the disappearance of the pre-announcement drift and date it to a single announcement. Over 2016–2023, the average cumulative return from the previous day’s open is flat and statistically insignificant before the announcement. Although the standard 24-hour drift window of [Lucca and Moench \(2015\)](#) still earns 0.21%, the gain is concentrated in the overnight gap, 0.19%, the only significant component; neither trading session contributes. The pre-announcement VIX decline documented alongside the drift disappears as well. For each announcement from 2012 to 2023, I treat it as a candidate break date and compare the average drift over the preceding 12 announcements with that over the 12 announcements beginning with the candidate. The largest downward shift lands on March 2016, significant at the 1% level. That announcement is the first after Cboe Global Markets (Cboe) listed the Wednesday SPX expirations on February 23, 2016, putting a same-day expiration on

the weekday of FOMC announcements. No candidate whose estimation window excludes the listing produces a significant decline. Even across the nearest four announcements on either side of the listing, the average drift falls from 1.05% to an insignificant 0.04%. The mechanism of the expiration introduction aligns with the model in [Hu et al. \(2022\)](#). They attribute the pre-announcement drift to the compensation for bearing the *uncertainty of the event uncertainty*, the variance of the event-impact variance, which resolves before the announcement. Once the Wednesday-expiring 0DTE options make the event uncertainty observable and tradable in real time, the higher-order uncertainty has little room to build, and thus the drift disappears.

Third, I show that the same accumulation-and-release pattern also appears, in attenuated form, around the release of the FOMC Minutes, which disclose the deliberations of the prior meeting. Because the Minutes carry less new information than the original announcement, the uncertainty buildup is lower, peaking at 0.34%, less than half the announcement peak, and the predictive structure is correspondingly noisier. The dominant pattern still matches FOMC announcement days. The implied event uncertainty in the run-up to the release significantly predicts realized volatility at the release window, with predictive power concentrated in the period immediately following the Minutes release. That a lower-information event reproduces the same dominant accumulation-and-release structure, scaled to its information content, supports the interpretation that event uncertainty accumulates rather than resolves before scheduled information deliveries.

My final finding extends the same predictive structure beyond scheduled news. Even on regular trading days, large price moves are not uniformly distributed across the session. The closing window of 15:50–16:00 ET concentrates more discontinuous volatility than any other ten-minute interval of the trading day, about 1.9 times the within-day average.¹ The closing window thus functions as a focal event, a regularly anticipated burst of price discovery, and the 0DTE option market prices it in the same accumulating fashion. Without a scheduled event, shocks arrive throughout the session rather than at a known instant, and discontinuous price moves likewise arrive at a finite rate. The option then implies a different pair of

¹At 15:50 ET, the market enters the closing-auction information phase: exchanges begin publishing indicative clearing-price/imbalance signals for the close at high frequency (NYSE: every second; Nasdaq: every 10 seconds until 15:55 ET, then every second). On NYSE, 15:50 ET is also the Market on Close (MOC)/Limit on Close (LOC) entry deadline; after this time, MOC/LOC cancellations are permitted only to correct bona fide errors and are rejected after 15:58.

quantities: the expected magnitude of discontinuous moves and their expected net direction. But they measure the same economic object—the market’s uncertainty about the size and tilt of the anticipated repricing. The expected magnitude predicts the closing-window realized discontinuous volatility from any sampling point in the day, with R^2 between 41% and 46%, while the expected net direction again has no predictive power. Within each segment of the day, the newly accumulated magnitude forecasts the closing burst, controlling for the initial expected magnitude level, ruling out both a rising risk premium and an artifact of the shrinking maturity in the measure’s denominator.

Taken together, all my findings describe a common mechanism: when the trading session contains an anticipated concentration of price discovery, whether driven by Fed communication or by closing-auction microstructure, option-implied event uncertainty accumulates in advance and resolves at the event, not before it.

Unlike the earlier option-based evidence, I use 0DTE options as the instrument for measuring this accumulation in real time. An option’s price reflects all of the risk between the moment of pricing and its expiration. The 0DTE option, expiring the same trading day, has a residual life measured in hours, and any imminent event is easier to separate from the background volatility. As expiration approaches, the residual life of the option contains fewer competing sources of risk, and the gap between at-the-money (ATM) implied volatility and spot volatility becomes increasingly informative about an impending discrete event rather than about everything that might happen over a distant expiration horizon. However, the long-maturity option blends the imminent event with every other source of uncertainty over its remaining life, weeks of background volatility, additional macroeconomic releases, and idiosyncratic news, so the imminent event is one risk among many. This cost of maturity is validated in the data. When I use options up to 5DTE, the predictive content of the announcement fades exactly as its shrinking share of the option’s total variance implies, while the closing-auction content survives, because every added day of an option’s life contains another close instead of just one.

With this measurement in hand, I make several contributions to the literature. Most directly, I add to the literature on FOMC announcements and asset pricing. I provide evidence that the pre-announcement window is where uncertainty builds rather than resolves. Predictive regressions show that the implied event uncertainty forecasts realized volatility only when the target window contains the announcement, indicating that the buildup carries genuine forecast

content about realized event outcomes rather than reflecting only a risk-premium pattern. I also document that the majority of the event variance resolves at the statement instead of the press conference. I further contribute to the literature on the pre-announcement drift. I show that the drift disappears in recent years, trace its disappearance to a single announcement rather than a gradual decline, and link the timing to the introduction of Wednesday-expiring 0DTE options, which complete the market for higher-order announcement uncertainty and thus eliminate the compensation for bearing it. In addition, I also contribute to the growing literature on ultra-short-maturity (USM) options. Existing empirical work on the discrete event component of 0DTE prices specifies the distribution and arrival process of price discontinuities at estimation, fitting fully parametric jump-diffusion models. The framework in this paper drops both restrictions and treats the two ways discontinuity risk can arrive, at a scheduled time or at a finite rate, within one semimartingale setting. For scheduled events, a cumulant expansion of the ATM pricing functional identifies the expected event volatility from the level of ATM implied volatility and the expected event skewness from its moneyness slope, under a near-Gaussian restriction that bounds how far the event distribution may depart from a Gaussian move without imposing a parametric family. For days without a scheduled event, I extend the short-maturity expansion of [Medvedev and Scaillet \(2007\)](#) from compound Poisson jumps to a compensator with time- and state-dependent intensity and an unspecified jump-size distribution; the ATM level then identifies the expected jump magnitude and the slope identifies the net direction, complementing the directional jump risk mapping of [Yan \(2011\)](#) under the Merton model. Finally, I extend the analysis beyond scheduled news. The accumulation-and-release mechanism documented for the FOMC announcement also operates around the closing auction, even though the close is not a news event in the conventional sense but a microstructural burst of price discovery. This connects two literatures that have largely developed in parallel: option-implied event-risk measurement, which has focused on scheduled macroeconomic releases, and the microstructure of closing-auction price discovery. The same accumulation pattern arises whenever there is a regularly anticipated concentration of price discovery, regardless of whether the driver is macroeconomic or microstructural.

Related Literature [Lucca and Moench \(2015\)](#) document that U.S. equity returns realize a disproportionate share of their daily expected return in the window before scheduled FOMC announcements, at unusually low realized volatility, a pattern that has come to be known as

the pre-FOMC announcement drift. [Ai and Bansal \(2018\)](#) establish that the existence of an announcement premium requires generalized risk-sensitive preferences as a foundation, and [Wachter and Zhu \(2022\)](#) model the drift as compensation for learning about a latent disaster probability that resolves at the announcement. [Cieslak, Morse, and Vissing-Jorgensen \(2019\)](#) attribute the drift to leakage of the policy stance through informal Fed communication. Such leakage would imply early revelation of the decision’s content—a partial resolution of event uncertainty before the announcement and a directional signal in option prices—neither of which I find from 0DTE options. Other studies explain the drift through additional risk channels beyond aggregate event uncertainty. [Ai, Bansal, and Han \(2021\)](#) and [Cocoma \(2025\)](#) build models in which disagreement between investors narrows as the announcement approaches, while [Laarits \(2025\)](#) traces the resolution to gradual learning about how the policy will be interpreted.

A parallel line of work documents the disappearing trend of the pre-announcement drift. [Boguth, Grégoire, and Martineau \(2019\)](#) find that over 2011–2017 the drift survives only ahead of announcements followed by a press conference, and [Kurov, Wolfe, and Gilbert \(2021\)](#) extend the sample and find it essentially gone over 2016–2019, with or without a press conference. I further sharpen the disappearance timing from a multi-year bracket to a single announcement and link it to the introduction of Wednesday-expiring 0DTE options.

The strand closest to this paper uses option markets to measure pre-announcement event risk directly, and it is here that my finding is in closest dialogue with prior work. [Beckmeyer, Branger, and Grünthaler \(2021\)](#) document that a left-tail premium builds up over the days before the announcement. I add to their work by extracting two complementary measures of the event impact, its uncertainty and its skewness, which together provide a richer characterization than a single one-sided premium. [Hu et al. \(2022\)](#) propose a heightened-uncertainty premium that compensates investors for bearing a specific kind of risk that resolves before the announcement and document a pre-FOMC VIX decline as supporting evidence. The relationship between their evidence and mine merits careful attention because the two are sometimes read as documenting the same economic content. The risk [Hu et al. \(2022\)](#) identify as resolving before the announcement is the *variance of the event-impact variance*: investors are uncertain about how large the announcement’s variance will be, and this higher-order uncertainty resolves before disclosure. On the other hand, the event uncertainty in this paper measures the volatility of the market’s reaction *itself*. Its rise before

the announcement therefore means the expected level of that variance is climbing. A declining variance of variance and a rising level of variance are not mutually exclusive, so my evidence is complementary to [Hu et al. \(2022\)](#) rather than in conflict. The pre-announcement VIX decline they document does not conflict with my reading either. VIX is a 30-day forward variance in which any single event mixes with weeks of higher moments and background variance, as my predictive tests across maturities show that the event’s forecast content fades sharply with added days to expiration. So VIX can decline through their higher-moment resolution channel while the event-specific variance is still building. Moreover, the disappearance of both the pre-announcement drift and the VIX decline is consistent with their model, in which the drift compensates the uncertainty of the event uncertainty and the VIX decline marks its resolution. Once that higher-order uncertainty can no longer build, neither should survive. [Johannes, Kaeck, and Seeger \(2026\)](#) structurally calibrate an FOMC event-volatility parameter from the cross-section of SPX options with maturity up to 60 days, find no average pre-announcement trend in that parameter, and assign the majority of the event variance to the press conference. I show different patterns, grounded in the different maturities of the option instruments. Their surface-wide parameter, estimated under a common background state, blends the imminent event with weeks of lower background risk, pulling the estimated background state toward those calmer periods. Before the announcement, this averages away the buildup in event uncertainty. After the statement, the same background estimate understates the elevated late-afternoon background volatility, causing the residual variation to be attributed to unresolved event uncertainty and thus inflating the press-conference share.

While the previous comparisons use ODTE options as a measurement instrument, a growing literature studies these options as objects of analysis in their own right, examining their pricing ([Bandi, Fusari, and Renò, 2024](#); [Almeida, Freire, and Hizmeri, 2024](#)), retail trading patterns ([Beckmeyer, Branger, and Gayda, 2023](#)), risk-factor profile ([Dorion, Orłowski, and Song, 2026](#)), and impact on intraday market volatility through market-maker hedging ([Brogaard, Han, and Won, 2023](#); [Adams et al., 2025](#); [Vasquez et al., 2025](#)).

Finally, my closing-auction findings connect to the microstructure of end-of-day price discovery. The closing auction concentrates volume, liquidity, and informational price formation in the final minutes of the trading day ([Cushing and Madhavan, 2000](#); [Jegadeesh and Wu, 2022](#); [Bogousslavsky and Muravyev, 2023](#)). The option market prices this microstructural focal event in the same accumulating fashion as scheduled macroeconomic releases, suggesting

that the accumulation pattern reflects how options price any anticipated concentration of price discovery, regardless of whether the trigger is informational or microstructural.

2. Identifying the Expected Price Discontinuity

This section describes the theoretical framework for extracting the expected price discontinuity uncertainty and asymmetry from options, and derives the framework under the two scenarios of days with and days without a scheduled event.

Under $\mathbb{Q}$, I consider the underlying price, S_t , and its log-return, x_t , which follow the Itô semimartingale:

$$\begin{aligned}\frac{dS_t}{S_{t-}} &= \sigma_t dW_t + \int_{\mathbb{R}} (e^x - 1) \tilde{\mu}(dt, dx), \\ dV_t &= d\sigma_t^2 = \kappa_v(\theta_v - \sigma_t^2) dt + \sigma_v \sigma_t dW_t^v,\end{aligned}\tag{1}$$

where W is a $\mathbb{Q}$ -Brownian motion with $\text{corr}(dW_t, dW_t^v) = \rho dt$, $\mu(dt, dx)$ is an integer-valued random measure with predictable compensator $\nu(dt, dx)$, and $\tilde{\mu}(dt, dx) := \mu(dt, dx) - \nu(dt, dx)$ is the martingale finite jump measure. In plain terms, prices move through a continuous component with stochastic volatility σ_t and through occasional discontinuities whose expected frequency and size are encoded in ν . I set the risk-free rate to zero throughout; all results extend to $r > 0$ by replacing S with the forward price.

The framework treats two cases, distinguished by whether the market knows in advance when an event will happen. On a day with a scheduled event, ν carries an atom at the known happening time; on a day without one, ν is absolutely continuous in time, so price jumps are expected to arrive at a finite rate. Each restriction governs how Black–Scholes–Merton (BSM) ATM implied volatility behaves relative to spot volatility as expiration approaches.

2.1 Days with a Scheduled Event

A scheduled event at the known time T_A delivers a single price move of random size Z_{T_A} . The compensator $\nu(dt, dx)$ then includes a fixed-time atom,

$$\nu(dt, dx) = \nu_t(dx) dt + \Pi_{A,t}(dx) \delta_{T_A}(dt), \quad \Pi_{A,t}(dx) := \mathcal{L}^{\mathbb{Q}}(Z_{T_A} \mid \mathcal{F}_t),\tag{2}$$

where $\delta_{T_A}(dt)$ is the Dirac measure placing unit mass at the known time T_A , so the event contributes exactly one price move at T_A , and $\Pi_{A,t}$ is the time- t conditional $\mathbb{Q}$ -law of the scheduled move: the market's expected distribution of the event impact, updated continuously as information arrives. The $\mathbb{Q}$ -martingale property of the underlying price is preserved by a predictable drift offset at T_A , and Z_{T_A} is conditionally independent of the diffusion paths given $\mathcal{F}_t$. Let $c_{n,t}^A$ denote the n th conditional cumulant of $\Pi_{A,t}$, write $\sigma_{A,t}^2 = c_{2,t}^A$ for the event variance, and standardize the higher cumulants as $s_{n,t}^A = c_{n,t}^A / \sigma_{A,t}^n$, so that $\sigma_{A,t}$ is the event volatility and $s_{3,t}^A$ the event skewness. Define the normalized moneyness $m = \log(F_t/K)/(\sigma_t\sqrt{T})$, where T is the annualized maturity, F_t is the forward, and K the strike.

Instead of assuming a pure Gaussian event move, I impose a near-Gaussian restriction on the shape of $\Pi_{A,t}$, which allows for small departures from Gaussianity.

ASSUMPTION 1 (Near-Gaussian event move): *There exists a perturbation parameter $\varepsilon \downarrow 0$ such that the standardized cumulants of the expected move satisfy*

$$s_{3,t}^A = O(\varepsilon), \quad s_{4,t}^A = O(\varepsilon^2), \quad s_{n,t}^A = O(\varepsilon^{n-2}), \quad n \geq 5, \quad (3)$$

uniformly in t . The parameter ε governs the departure of $\Pi_{A,t}$ from the Gaussian benchmark and is distinct from the maturity T .

The rates in Assumption 1 mirror the standard cumulant ordering underlying an Edgeworth expansion (e.g., Hall, 2013). At $\varepsilon = 0$ the assumption reduces to the Gaussian announcement move. Consider an announcement that lies between now and the expiration of options. That is, for $t < T_A < t + T$ and $T \downarrow 0$, the conditional independence of the event move and the background components makes their variances add, which yields the ATM implied variance decomposition in the following proposition.

PROPOSITION 1 (Days with a scheduled event): *Let the compensator be (2) with $t < T_A < t + T$, let Z_{T_A} be conditionally independent of the diffusion and non-scheduled jump components given $\mathcal{F}_t$, and let Assumption 1 hold with $\sigma_{A,t}$ bounded away from zero and $\int_{\mathbb{R}} |x|^3 \nu_t(dx) < \infty$. Then, as $T \downarrow 0$ and $\varepsilon \downarrow 0$,*

$$I(0, T; \sigma_t)^2 T \cong \sigma_{A,t}^2 + \sigma_{N,t}^2 T + \sigma_t^2 T + O(\varepsilon \sigma_{A,t}^3, \varepsilon^2 \sigma_{A,t}^2, \dots) + O(T), \quad \sigma_{N,t}^2 = \int_{\mathbb{R}} x^2 \nu_t(dx), \quad (4)$$

where $I(0, T; \sigma_t)$ is the BSM ATM implied volatility at maturity T with normalized moneyness $m = 0$, $\sigma_{N,t}^2$ is the variance contribution of non-scheduled discontinuities, and the final $O(T)$ term collects the remaining order- T contributions from third- and higher-order cumulants of the background components. The slope of the ATM implied volatility smile is

$$\left. \frac{\partial I}{\partial m} \right|_{m=0} \cong -\frac{\sigma_t}{6} s_{3,t}^A + O(\varepsilon^2) + O(T). \quad (5)$$

The proof is in [Appendix A](#).

The expansion (4) carries the higher-order terms of the event variance and cumulants, which are suppressed because $\varepsilon \downarrow 0$, and in practice the maximum value of expected event volatility, $\sigma_{A,t}$, is about 1%. So, the higher-order terms are small relative to the leading term. The expected event volatility can be estimated by

$$\mathcal{U}_t^A = \sqrt{(I(0, T; \sigma_t)^2 - \sigma_t^2 - \sigma_{N,t}^2)T}, \quad \mathcal{U}_t^A \approx \sigma_{A,t}. \quad (6)$$

The level is silent on the direction of the expected move; the asymmetry is encoded in the slope, and the expected event skewness can be estimated by

$$\mathcal{S}_t^A = -\frac{6}{\sigma_t} \left. \frac{\partial I}{\partial m} \right|_{m=0}, \quad \mathcal{S}_t^A \approx s_{3,t}^A. \quad (7)$$

The ATM slope identifies the asymmetry of jump risk in both day types, though the quantity differs; under a continuous compensator it recovers the net jump direction, as in [Yan \(2011\)](#), the case of [Section 2.2](#).

2.2 Days without a Scheduled Event

Without a scheduled event, price discontinuities still occur, but their expected number scales with the length of the interval rather than concentrating at a known time. The compensator is absolutely continuous in time,

$$\nu(dt, dx) = \nu_t(dx) dt, \quad (8)$$

satisfying $\int_{\mathbb{R}}(e^x - 1 - x) \nu_t(dx) < \infty$. In the absence of the scheduled event term, the event volatility $\sigma_{A,t}$ is essentially zero, and every term in equation (4) is thus $O(T)$. Dividing by T gives $I(0, T; \sigma_t)^2 = \sigma_t^2 + \sigma_{N,t}^2 + O(1)$, where the remainder collecting the higher-order cumulants is of the same order as $\sigma_{N,t}^2$ and does not vanish as $T \downarrow 0$, which makes $\sigma_{N,t}$ unidentifiable from the level of ATM implied volatility.

However, the level still identifies a different variable with the same economic interpretation: the *expected jump magnitude*, $\int_{\mathbb{R}} |e^x - 1| \nu_t(dx)$, which enters the smile at order $\sqrt{T}$. It co-moves with $\sigma_{N,t}$ and therefore measures the same economic object, the size of expected jumps.

PROPOSITION 2 (Days without a scheduled event): *Let the compensator be (8) with finite activity, $\nu_t(\mathbb{R}) < \infty$. As $T \downarrow 0$,*

$$I(m, T; \sigma_t) = \sigma_t + \left(-\frac{\rho\sigma_v}{4}m - \kappa_J(\nu_t) \frac{N(m)}{n(m)} + \kappa_{J+}(\nu_t) \frac{1}{n(m)} \right) \sqrt{T} + O(T), \quad (9)$$

where $N(\cdot)$ and $n(\cdot)$ are the standard normal distribution and density functions, $\kappa_J(\nu_t) = \int_{\mathbb{R}}(e^x - 1) \nu_t(dx)$, $\kappa_{J+}(\nu_t) = \int_{\mathbb{R}_+}(e^x - 1) \nu_t(dx) > 0$, and $\kappa_{J-}(\nu_t) = \int_{\mathbb{R}_-}(e^x - 1) \nu_t(dx) < 0$. At the money, the level and the slope of the smile satisfy

$$I(0, T; \sigma_t) = \sigma_t + \sqrt{\frac{\pi}{2}} (\kappa_{J+}(\nu_t) - \kappa_{J-}(\nu_t)) \sqrt{T} + O(T) \quad (10)$$

and

$$\left. \frac{\partial I}{\partial m} \right|_{m=0} = - \left(\frac{\rho\sigma_v}{4} + \kappa_J(\nu_t) \right) \sqrt{T} + O(T). \quad (11)$$

The proof is in Appendix B.

The expansion (9) extends the short-maturity expansion of Medvedev and Scaillet (2007), who derive it for compound-Poisson jumps with constant intensity, to a compensator with time- and state-dependent intensity and an unspecified jump-size distribution. The expected magnitude of discontinuous moves, $\kappa_{J+} - \kappa_{J-}$, is therefore the leading term in the wedge between ATM implied volatility and spot volatility and can be estimated by

$$\mathcal{M}_t^{N,J} = \sqrt{\frac{2}{\pi}} \frac{I(0, T; \sigma_t) - \sigma_t}{\sqrt{T}}, \quad \mathcal{M}_t^{N,J} \approx \kappa_{J+}(\nu_t) - \kappa_{J-}(\nu_t). \quad (12)$$

As in the scheduled-event case, the slope identifies the asymmetry of jump risk, but now the net directional jump κ_J is the relevant quantity: the expected net jumps can be estimated by inverting the slope equation (11),

$$\mathcal{J}_t^N = -\frac{1}{\sqrt{T}} \left. \frac{\partial I}{\partial m} \right|_{m=0} - \frac{\rho\sigma_v}{4}, \quad \mathcal{J}_t^N \approx \kappa_J(\nu_t). \quad (13)$$

The mapping from the smile slope at ATM to directional jump risk parallels Yan (2011), who obtains it under the Merton model. The slope (11) alone, however, identifies only the net jumps $\kappa_{J+} + \kappa_{J-}$; the level (10) contributes the total magnitude $\kappa_{J+} - \kappa_{J-}$, and combining them gives the signed components of expected jumps:

$$\begin{aligned} \mathcal{J}_t^{N+} &= \frac{1}{2}(\mathcal{J}_t^N + \mathcal{M}_t^{N,J}), & \mathcal{J}_t^{N+} &\approx \kappa_{J+}(\nu_t), \\ \mathcal{J}_t^{N-} &= \frac{1}{2}(\mathcal{J}_t^N - \mathcal{M}_t^{N,J}), & \mathcal{J}_t^{N-} &\approx \kappa_{J-}(\nu_t). \end{aligned} \quad (14)$$

This recovers the expected positive and negative jumps separately.

3. Data

I use S&P 500 Index (SPX) options from September 2016 to December 2023, drawing on the Cboe Global Markets (Cboe) Option Intervals dataset, which records option quotes at the one-minute frequency alongside a synchronized level of the underlying index. For a liquid and tradable proxy of the underlying price, I rely on millisecond-stamped E-mini S&P 500 (ES) futures trades from Tickdata Inc., and I take risk-free rates from the OptionMetrics zero-coupon yield curve, held fixed within each trading day. Each option price is the midpoint of its one-minute quote. I retain PM-settled contracts with at most 5 business days to expiration (DTE). To limit microstructure noise from the thin quoting near the open and close, I restrict the option sample to the 10:00 to 15:40 ET window, while keeping the full 9:30 to 16:00 ET session for the ES futures used to estimate realized volatility.

3.1 Stochastic Maturity

Calendar time mismeasures the remaining risk of an option whose life is a few hours: variance flows unevenly through the trading day and almost stops overnight and on weekends, so two calendar hours at the open and two at midday contain different amounts of risk. [Bondarenko \(2026\)](#) proposes measuring the maturity of a USM option in units of variance accumulation rather than calendar time or only trading hours. Because option prices identify total variance, not volatility and maturity separately, a clock that misstates how variance accrues distorts every subsequent implied volatility calculation. The variance clock removes much of the resulting irregularity at the short end of the surface. Building on [Bondarenko \(2026\)](#), I adopt this variance clock in the minute-level implementation. The weight of each minute of day s is $\lambda(s) = \bar{\sigma}_s^2 / \bar{\sigma}_{\text{RTH}}^2$, where $\bar{\sigma}_s^2$ is the average one-minute realized variance at s and $\bar{\sigma}_{\text{RTH}}^2$ is the average one-minute realized variance across all regular trading hours (RTH) minutes.² The normalization sets the average RTH minute to one. The maturity of an option observed at t and expiring at time $\mathcal{T}$ is the accumulated weight between the two,

$$T = \frac{1}{\Omega} \sum_{s \in (t, \mathcal{T}]} \lambda(s), \quad (15)$$

where $\Omega \approx 181,750$ is the number of RTH-equivalent minutes in a year.³ All annualized magnitudes and maturities in this paper use this clock.

3.2 ATM Implied Volatility

For each expiration, I first recover the forward price by selecting the pair at the strike, $K_{t,T}^*$, whose call and put prices are closest, so the extraction uses the pair nearest the money, where quote errors in the price difference are smallest. The forward is then computed from

²As in [Dorion, Orlowski, and Song \(2026\)](#), one calendar week is divided into six regimes: Global Trading Hours (GTH, 18:00 to 09:30 ET), Regular Trading Hours (RTH, 09:30 to 16:00 ET), Post-market (Post, 16:00–16:15 and 16:30–17:00 ET), CME settlement halt (HALT, 16:15–16:30 ET), the technical break (BRK, Monday to Thursday, 17:00–18:00 ET), and weekends and holidays (HOL, weekends start from 17:00 ET on Friday).

³A calendar year contains 252 trading days, each with a GTH, RTH, HALT, and Post session. Of these trading days, 191 days have a standard 60-minute BRK period, and 61 days include pre-holiday gaps (17:00–18:00 ET of 52 Fridays and 9 pre-holiday days) classified as HOL. The remaining 113 days are full 24-hour HOL.

put-call parity on this pair, $F_{t,T} = K_{t,T}^* + e^{rT}(C_{t,T}(K_{t,T}^*) - P_{t,T}(K_{t,T}^*))$. Given the forward, I define the ATM pair as the call with the smallest strike above $F_{t,T}$ and the put with the largest strike below it; these ATM options are the main instruments in the paper. The ATM implied volatility is the BSM implied volatility at the forward, obtained by interpolating the implied volatilities of the ATM put and call,

$$\hat{I}_t(0, T) = \frac{K_{C,ATM} - F_{t,T}}{K_{C,ATM} - K_{P,ATM}} \hat{I}_t(0, T; P) + \frac{F_{t,T} - K_{P,ATM}}{K_{C,ATM} - K_{P,ATM}} \hat{I}_t(0, T; C), \quad (16)$$

where $K_{P,ATM}$ and $K_{C,ATM}$ are the ATM put and call strikes and $\hat{I}_t(0, T; P)$ and $\hat{I}_t(0, T; C)$ are their implied volatilities.

4. Measurement

The framework of Section 2 defines the event measures through quantities that must be estimated from the data. This section explains how I compute each one. The estimates draw on two sources, the high-frequency path of ES futures and the panel of option prices, and build on one another. The continuous component of the realized volatility of the futures is used to tune the option-based spot volatility estimator, and the resulting spot volatility is an input to the option-based measures of the expected discontinuous price movement. The section proceeds as follows. Section 4.1 defines the realized volatility measures, Section 4.2 introduces the option-based spot volatility estimator, and Section 4.3 presents the implied event measures and their non-scheduled counterparts.

4.1 Realized Volatility

To estimate the realized volatility of the underlying, I use five-second log returns, Δx , of ES futures. For an intraday window $[t, t+h)$ containing n log returns $\Delta_i x$, realized volatility is

$$RV_{t,t+h} = \sqrt{\sum_{i=1}^n |\Delta_i x|^2}, \quad (17)$$

which estimates the volatility of all price movement over the window. [Bollerslev and Todorov \(2011\)](#) show that the price's quadratic variation separates into a continuous part, accumulated

from the many small moves of the diffusion, and a discontinuous part, carried by the few large moves. To recover the continuous part, I use the truncated variation as in [Bollerslev and Todorov \(2011\)](#) by keeping only returns below a threshold:

$$\text{CV}_{t,t+h}^2 = \sum_{i=1}^n |\Delta_i x|^2 \mathbf{1}(|\Delta_i x| \leq \alpha_t n^{-0.49}). \quad (18)$$

Following [Bollerslev, Todorov, and Zengzhi Li \(2012\)](#), I set the threshold as $\alpha_t = \tau \sqrt{\text{RV}_{t-w,t}^2 \wedge \text{BV}_{t-w,t}^2}$, which scales from the preceding w minutes of realized variance and bipower variation, $\text{BV}_{t-w,t}^2 = \frac{\pi}{2} \sum_i |\Delta_{i-1} x| |\Delta_i x|$. I select $\tau = 2.5$ and $w = 30$ minutes, and [Figure A.I](#) in the Appendix shows the intraday patterns of CV under alternative values of τ and w . The standard practice scales the previous day's variance by a time-of-day factor that captures the average intraday U-shape, whereas I use the variance from a trailing intraday window, which adapts to the prevailing volatility level and embeds the intraday patterns automatically. Moreover, on event days the unconditional U-shape is the wrong benchmark. The pre-announcement hours are calmer than the average and the post-release hour, so a fixed U-shape correction truncates too aggressively in the morning and too loosely after the release. The discontinuous component of realized volatility is the remainder,

$$\text{JV}_{t,t+h} = \sqrt{(\text{RV}_{t,t+h}^2 - \text{CV}_{t,t+h}^2)_+}, \quad (19)$$

where $(x)_+ = \max(x, 0)$. The difference between the realized and continuous variance measures the contribution of discontinuous moves to total variance.

4.2 Spot Volatility from Options

The realized CV is not by itself an adequate spot volatility input for [Section 2](#). It is estimated from a trailing window, so CV lags the true volatility when it moves, and the annualization that makes a futures-based measure comparable to an option-implied one is itself uncertain. Option prices provide a risk-neutral spot measure that avoids all the problems. I use CV only as a tuning input, where these weaknesses are second order, and use options with the shortest available maturity at each timestamp.

Todorov (2019) proposes a spot volatility estimator to recover the instantaneous diffusive volatility from the USM options. Following Todorov (2019) and Chong and Todorov (2025), the conditional characteristic function of the log return is spanned by out-of-the-money (OTM) option prices,

$$\widehat{L}_t(u) = 1 - (u^2 + iu) \sum_{j=2}^{N_k} \frac{e^{iu \log(K_{j-1}/F_t)}}{K_{j-1}} O_t(K_{j-1}) \Delta \log K_j, \quad (20)$$

where $O_t(K)$ is the OTM option price at strike K and N_k is the number of strikes. The estimation of the spot volatility, σ_t , in Section 2 is

$$\text{SV}_t = \sqrt{-\frac{2}{\widehat{u}^2 T} \log |\widehat{L}_t(\widehat{u})|}, \quad \text{SV}_t = \widehat{\sigma}_t. \quad (21)$$

For small T , the quantity $-\frac{2}{u^2 T} \log |L_t(u)|$ decomposes into the spot variance σ_t^2 and a bias from the price discontinuities, which is bounded by $O_p(u^{r-2})$ with $r \leq 1$ for finite-variation jumps. A larger u shrinks this bias, but because $|L_t(u)|$ decays toward zero as u grows, it also amplifies the estimation noise. To balance the effects, Todorov (2019) shows the best choice is to hold $u^2 \sigma_t^2 T$ at a fixed constant. The evaluation point is therefore set as the inverse square root of the remaining diffusive variance,

$$\widehat{u}_t = \sqrt{\frac{-2 \log \ell}{\text{CV}_t^2 T}}, \quad \ell = 0.1, \quad (22)$$

which targets $|\widehat{L}_t(\widehat{u})| \approx \ell$ when the diffusive variance over the option's remaining life is close to $\text{CV}_t^2 T$.

Chong and Todorov (2025) implement the same estimator with an option-implied total variance in place of $\text{CV}_t^2 T$ in equation (22). On most days the two choices are interchangeable; on FOMC announcement days they diverge in opposite directions. The implied variance of 0DTE options is dominated by the expected event variance, which rises through the morning and collapses at the announcement, while the realized volatility shows the opposite pattern, as introduced in Section 5.1. The implied variance therefore moves up before the announcement, pushing $\widehat{u}_t$ down. The estimator then fails to filter out the expectation of event variance in the option prices. The realized CV^2 tracks the diffusive variance without

the event component, so it keeps $\hat{u}_t$ stable and unbiased by the event.

Using an estimator based on high-frequency futures return data as an input into an option-based estimator is internally consistent because spot volatility is a property of the price path: an equivalent change of measure alters the drift and the compensator but not σ_t , so the spot volatility should be the same under $\mathbb{P}$ and $\mathbb{Q}$, and [Chong and Todorov \(2025\)](#) find that spot volatility estimated from the two markets agrees closely in the data. Moreover, CV enters only through the evaluation point $\hat{u}$, not through the level of the estimator, so the bias in CV perturbs SV only at second order.

Using the shortest available maturity without adjustment does not produce a valid estimate at every minute, because too few OTM quotes may remain or equation (21) may fail to produce a finite value, especially as the expiration nears. To keep the maximum data, I start from the shortest available maturity and, when its estimate is invalid, fall back to the next longer maturity until a valid estimate is obtained. Appendix Figure A.II and Appendix Table A.I compare this fallback procedure with two alternatives. Dropping invalid minutes mostly removes late-afternoon observations, when estimation failures for the shortest maturity are most frequent, whereas the fallback procedure restores full coverage through the day. Averaging the two shortest valid maturities instead introduces a different bias. The longer maturity mixes the elevated late afternoon on FOMC announcement days with the calmer days that follow, so the two-maturity series underestimates the high spot volatility on the afternoon after 14:00 ET.

4.3 Implied Price Discontinuity

After estimating the spot volatility, I can estimate the expected price discontinuity under the two scenarios illustrated in Section 2.

On days with a scheduled release, using the ATM implied volatility and SV, I estimate the expected event volatility from equation (6) by

$$\hat{\mathcal{U}}_t^A = \sqrt{(\hat{I}_t(0, T)^2 - \text{SV}_t^2)_+ T}. \quad (23)$$

I do not net out the non-scheduled jump variance $\sigma_{N,t}^2$ for two reasons. First, the expected event variance is the leading term in the USM implied variance. Second, the pre-FOMC announcement period is usually quiet, with realized volatility even lower than on non-FOMC

days (Ai, Bansal, and Han, 2021). The non-scheduled jump variance is therefore negligible during the pre-announcement calm period. The slope of the smile at the ATM further recovers the expected event skewness from equation (5). The slope is estimated minute by minute from a fifth-order polynomial regression of the smile on the normalized moneyness, fitted to OTM quotes ($K_P < F_t < K_C$),

$$I(m, T) = \alpha + \sum_{k=1}^5 \hat{\gamma}_k m^k + \epsilon. \quad (24)$$

Polynomial smile fitting follows the smile-implied approach of François and Stentoft (2021), who fit the smile in simple moneyness. Using the slope of the smile at $m = 0$, the estimator of the expected event skewness is therefore

$$\hat{\mathcal{S}}_t^A = -\frac{6}{\text{SV}_t} \hat{\gamma}_{1,t}. \quad (25)$$

On days without a scheduled release, the magnitude and direction of non-scheduled discontinuity risk are also identified from the ATM implied volatility and slope. The estimator of the expected jump magnitude is

$$\widehat{\mathcal{M}}_t^{N,J} = \sqrt{\frac{2}{\pi}} \frac{\hat{I}_t(0, T) - \text{SV}_t}{\sqrt{T}}. \quad (26)$$

The expected net jump and its directional components are estimated from the ATM slope as

$$\hat{\mathcal{J}}_t^N = -\frac{\hat{\gamma}_{1,t}}{\sqrt{T}} - \frac{\widehat{\rho\sigma_v}}{4}, \quad \begin{cases} \hat{\mathcal{J}}_t^{N+} = \frac{\hat{\mathcal{J}}_t^N + \widehat{\mathcal{M}}_t^{N,J}}{2}, \\ \hat{\mathcal{J}}_t^{N-} = \frac{\hat{\mathcal{J}}_t^N - \widehat{\mathcal{M}}_t^{N,J}}{2}, \end{cases} \quad (27)$$

where the estimation of the leverage term $\widehat{\rho\sigma_v}$ follows Chong and Todorov (2024) and is detailed in Appendix C.

5. FOMC Scheduled Events

This section documents the option market's expectations of the FOMC's scheduled events, the policy announcements and the Minutes releases. The event-day analysis covers the hours

when the event still lies ahead and the high-fluctuation interval just after it in both cases: the 10:00–15:00 ET window on the FOMC announcement days and the 10:00–14:15 ET window on the Minutes release days. The remaining intraday windows are the scenarios without scheduled events, which are analyzed in Section 6.

5.1 FOMC Announcement

The FOMC statement arrives at 14:00 ET, followed on most announcement days by a press conference at 14:30 ET.⁴ The timing is known to the minute while the content is not. This subsection first documents how the 0DTE option market prices the incoming events through the day, then tests whether that pricing carries information about the realization.

5.1.1 Intraday Patterns of Event Uncertainty

Figure 1 shows the intraday patterns of volatility on FOMC announcement days. The top row plots the intraday path of one-minute annualized RV from equation (17), averaged across announcement days and split by the policy outcome. The morning is quiet. Realized volatility drifts down from roughly 15% to 10%, the intraday counterpart of the pre-announcement calm documented by [Lucca and Moench \(2015\)](#) and [Ai, Bansal, and Han \(2021\)](#). At 14:00 ET it spikes, reaching 95% on the average announcement day, 132% on days when the target rate changes, and 71% on days it holds, with a second burst, to 37% on average, once the press conference begins at 14:30 ET. The middle row separates realized volatility into its continuous component and the component generated by discontinuous price moves, estimated from equations (18) and (19). The spike loads almost entirely on the discontinuous component: on the average announcement day JV peaks at 94%, while the highest CV is below 29%. The announcement impact arrives as a single discrete repricing at a known time, the scheduled-event scenario of Section 2.1, rather than as an intensification of ordinary trading ([Ederington and Lee, 1993](#); [Fleming and Remolona, 1999](#); [Anderson et al., 2003](#); [Savor and Wilson, 2013](#)).

The bottom row contains the measurement from 0DTE options. While realized volatility declines through the morning, the 0DTE ATM implied volatility climbs from 22% to 32%, and the option-implied spot volatility SV, estimated from equation (21), drifts down from

⁴Before 2019, the press conference followed half of the meetings each year; since January 2019, every scheduled meeting is followed by a press conference.

15% to about 10%; the wedge between the two widens from 7% to 22% as the announcement approaches, and it is twice as wide on days that end with a rate change (34%) as on hold days (16%). Without a scheduled event, equation (10) shows the wedge enters only at order $\sqrt{T}$, so it shrinks to zero as the option expires. Here it does the opposite: a scheduled event adds a fixed event variance to the option’s total variance (equation (4)). The widening wedge is the signature of accumulating event uncertainty.

Equation (6), implemented as in Section 4.3, converts the wedge into the expected event uncertainty $\hat{\mathcal{U}}_t^A$, the standard deviation of the price move the market expects the announcement to cause. Furthermore, the slope of the smile at the ATM recovers the expected event skewness $\hat{\mathcal{S}}_t^A$, the directional tilt of that move’s distribution. Figure 2 plots both measures through the announcement day, the event uncertainty in the top row and the directional tilt in the bottom. The event uncertainty rises monotonically from 0.60% at 10:00 ET to 0.78% before the release of the announcement, a 30% accumulation over four hours with no visible retracement. The level value is consistent with Johannes, Kaeck, and Seeger (2026), who find average event volatility of 0.75% for the 2012–2023 period at one hour before the announcement. The level also discriminates across meetings. On days with a change in the target rate the morning path runs from 0.83% to 1.08%, nearly twice the level on hold days, so the option market anticipates which meetings carry large impacts. Resolution happens only after information *delivery*, and it happens in two steps. At 14:00 ET, the event uncertainty collapses by more than half, to 0.34%; it then rebuilds to 0.42% ahead of the press conference at 14:30 ET, a second scheduled disclosure that concentrates market attention (Boguth, Grégoire, and Martineau, 2019), and decays to 0.22% once the chair begins speaking. The rebuild is strongest on change days, where the event uncertainty climbs from 0.42% back to 0.59% in the 20 minutes before the press conference, the meetings whose delivery leaves the most to explain. The pre-announcement window is therefore where event uncertainty builds, peaking at the moment the statement crosses the wire; it is not where event uncertainty dissipates.

The distribution asymmetry reinforces the same picture, as it carries no advance information about the outcome. Throughout the morning $\hat{\mathcal{S}}_t^A$ is negative, -0.53 on the average announcement day. The market assigns more weight to a large downside realization than to a comparable upside one, consistent with Beckmeyer, Branger, and Grünthaler (2021), who study the option-implied distribution of the FOMC announcement move and document

a premium for left-tail insurance that builds over the days before the meeting. Moreover, the asymmetry is flat through the pre-announcement period. Between 10:00 and 14:00 ET it drifts by only 0.06 on the average announcement day, and only after the release does it climb toward -0.40 as the priced downside scenario expires. The level differs across meetings, from -0.39 on target rate change days to -0.61 on hold days, yet the within-day path is equally flat in every subsample. A leaked policy decision would shift the implied distribution toward the leaked direction as the release approaches. The informal policy communication (Cieslak, Morse, and Vissing-Jorgensen, 2019) and informed order flow before scheduled releases (Bernile, Hu, and Tang, 2016) can both carry directional information in advance. But the distribution implied by the option market never shifts toward either side. The market becomes steadily more uncertain about the announcement’s outcome while learning nothing about its sign.

5.1.2 Resolution after Statement versus Press Conference

The two-step resolution shows most of the event uncertainty is resolved by the statement, not the press conference. In variance units, the collapse at 14:00 ET, from 0.78% to the trough of 0.34%, removes 81% of the priced event variance within minutes of the release. Part of it returns as the market re-prices the press conference. The rebuilt level of 0.42% still puts the statement’s resolved share at 71% and leaves 29% to the press conference. The press conference accounts for the minority of the event risk, and it is stable across decisions, 70% resolved by the statement on target rate change days and 72% on hold days.

Johannes, Kaeck, and Seeger (2026) reach the opposite result. They re-estimate the remaining event risk ten minutes before the press conference from their parametric model and attribute 62% of the event variance to the press conference, and 72% in 2020–2023. However, the realized record points the other way. The discontinuous volatility, JV, spikes to 94% at the statement against 19% once the press conference begins. The spike sizes measure the resolution. The repricing at a disclosure is the market’s revision of its conditional expectation of fundamentals, and the law of total variance splits the outstanding event variance into the variance of that revision and the expected uncertainty left for later disclosures, the mechanism Ai, Han, and Xu (2026) formalize as information-driven volatility. So the larger spike at the statement identifies it as the disclosure that carried most of the information. The premium evidence of Johannes, Kaeck, and Seeger (2026) also agrees: straddle returns

spanning the statement are significantly negative, while straddles held through the press conference earn returns indistinguishable from zero, a pattern the authors flag as curious under a majority press-conference share and the direct prediction of a minority one. Under the two-step resolution in this paper, the pattern is not a puzzle. The premium is earned where the uncertainty resolves, at the statement, and the press conference carries too little remaining uncertainty to command detectable compensation.

Their opposite conclusion comes from how the remaining event risk is measured. The statement has lifted the background volatility, but their model re-fits the cross-section of options with up to 60 days to maturity with one common background state and one set of dynamics. The longer maturities push the common state to the calmer weeks for the life after the event. Although they add the periodicity adjustment for each maturity separately, it is the same average pattern on every day, estimated from prior years, and cannot supply an event afternoon’s surge and decay. The fitted background therefore falls short of the level the market is experiencing in that late afternoon, thus inflating the remaining event uncertainty share. The measurement from only the 0DTE options in this paper is immune to the problem. An option that expires at today’s close prices nothing beyond the event day, so the calmer weeks after the event day never enter the instrument.

5.1.3 Evidence of No Early Resolution

The accumulation pattern alone cannot distinguish an actual increase in the expected event uncertainty from a rising risk premium, as both produce the same upward path. Predictive regressions help separate the two interpretations, as a genuine expectation of a larger event predicts a larger realized move, while a premium need not. I sample $\hat{\mathcal{U}}^A$ and $\hat{\mathcal{S}}^A$ minute by minute over 10:00–13:00 ET and forecast the realized volatility and the log return of every five-minute window from 13:00 to 15:00 ET, the period from one hour before the announcement to the end of the press conference,

$$\text{RV}_{t,t+5\text{min}} = \alpha + \beta_{\mathcal{U}} \hat{\mathcal{U}}_s^A + \beta_{SV} \text{SV}_s + \varepsilon_{t,t+5\text{min}}, \quad (28)$$

$$r_{t,t+5\text{min}} = \alpha + \beta_{\mathcal{S}} \hat{\mathcal{S}}_s^A + \beta_{SV} \text{SV}_s + \varepsilon_{t,t+5\text{min}}, \quad (29)$$

where s is the time at which the regressors are estimated and $[t, t + 5\text{min})$ is the forecast window, with $s < t$. Each pair of sampling time and forecast window is a separate time-series

regression across announcement days, and the coefficients measure whether $\hat{U}_s^A$ and $\hat{S}_s^A$ forecast the size and direction of the afternoon move beyond what is already predicted by the SV control. Figure 3 reports the results. The forecast content for the realized volatility is concentrated precisely at the event time. For the forecast windows ending *before* 14:00 ET, the coefficient of $\hat{U}_s^A$ is *not* statistically significant at nearly every sampling time s , and the moderate explanatory power (adjusted R^2 of 29%–53% at $s = 13:00$ ET) in regression (28) comes from the spot volatility control, the standard persistence channel of ordinary volatility forecasting. Once the window contains the announcement, the R^2 jumps to 78%–82%, and the coefficient of $\hat{U}_s^A$ is significant at the 1% level for every sampling point s from 10:00 to 13:00 ET. The R^2 peaks again at 85% for the window including the start of the press conference, in line with the second spike in realized volatility. The expectation read from option prices hours before the release explains the bulk of the cross-day variation in the size of the announcement move; the accumulating event uncertainty carries genuine forecasting power about the realization rather than reflecting only compensation for bearing announcement risk. More telling is the absence of any forecast content before the announcement. In most forecast windows ending before 14:00 ET, $\hat{U}_s^A$ cannot predict realized volatility, so the uncertainty does not resolve gradually through the morning and is resolved only when the statement is released.

Direction remains unpredictable. The bottom panel of Figure 3 reports the predictive regression of equation (29), the log return regressed on the directional asymmetry. The R^2 hovers near zero at most forecast windows, with a single exception (35%) in the window ending 14:45 ET, during the press conference, and the coefficient of $\hat{S}_s^A$ is statistically insignificant at almost all pairs. The asymmetry is further clean evidence that nothing resolves early. The option market forecasts how large the move will be but never learns which way it will go, and the directional distribution has little predictive power and stays flat up to 14:00 ET instead of shifting toward the outcome, the signature of a decision whose content is withheld until it is delivered.

Alternative method. The accumulation of event uncertainty could in principle inherit error from the spot volatility input. I therefore re-estimate the event uncertainty without any spot volatility input, using only the term structure of ATM implied volatility across the 0DTE and 2DTE options. The event variance enters the ATM implied variance at rate $1/T$

at every maturity, while the diffusive and background jump variances enter at the same level regardless of maturity, so differencing the two maturities isolates $\sigma_{A,t}$, in the same spirit as the term-structure identification of earnings announcement risk in [Dubinsky et al. \(2019\)](#). The alternative event uncertainty is estimated as

$$\hat{\mathcal{U}}_t^{A,\text{alt}} = \sqrt{\frac{\hat{I}_t(0, T_1)^2 - \hat{I}_t(0, T_2)^2}{1/T_1 - 1/T_2}}. \quad (30)$$

Figure 4 shows the alternative measure reproduces both findings. It climbs from 0.63% to 0.73% before the announcement and collapses at 14:00 ET. The predictive result shows the same structure, with adjusted R^2 of 68%–85% for windows containing or following the release, while the R^2 is only 29%–53% for windows ending before 14:00 ET, and the coefficient of $\hat{\mathcal{U}}_s^{A,\text{alt}}$ is not statistically significant at most sampling times s . Neither the accumulation nor the predictive power is an artifact of how spot volatility is measured.

5.1.4 Expectation versus Risk Premium

The predictive regression (28) shows that the expected event uncertainty tracks the realized size of the announcement move, but it does not establish whether the pre-announcement accumulation contributes to that tracking. All the forecast content in Figure 3 could sit in the 10:00 ET level. If the expectation of the announcement impact were fully priced before the morning began, a rising risk premium added to that level would let $\hat{\mathcal{U}}_s^A$ forecast the realized move at every later sampling time. The forecast content would then be inherited from the morning open, not accumulated through the pre-announcement hours.

Table I reports the statistics of the intraday accumulation over the pre-announcement morning. The accumulation is measured by the changes in the expected event uncertainty over the two halves of the morning, 10:00–12:00 ET and 12:00–13:59 ET, together with a fitted trend. The fitted trend is the average slope from regressing each day’s minute-level $\hat{\mathcal{U}}^A$ path on time over 10:00–13:59 ET. Both increments are positive, 0.10% and 0.03%, the fitted trend reports an increase of 0.038% per hour, and all three are statistically significant at the 5% level or better. The trend survives splitting by policy outcome, 0.050% per hour on days the target rate changes and 0.031% on hold days, both significant at the 1% level, so the climb is not driven by a single type of meeting.

To further examine the contribution of the pre-announcement accumulation, I apply the following predictive regression to forecast the realized discontinuous volatility,

$$JV_{t,t+h} = \alpha + \beta_{\Delta U} \Delta \hat{U}_s^A + \beta_U \hat{U}_{\text{ref}}^A + \beta_{SV} SV_s + \varepsilon_{t,t+h}, \quad (31)$$

where $\Delta \hat{U}_s^A = \hat{U}_s^A - \hat{U}_{\text{ref}}^A$ is the event uncertainty accumulated between the reference time and the sampling time s , and the forecast window $[t, t+h)$ is fixed at the announcement window [14:00, 14:15) ET or the press conference window [14:30, 15:00) ET. Each increment enters a separate regression across announcement days, the first from a 10:00 ET reference to $s = 12:00$ ET and the second from 12:00 ET to $s = 13:59$ ET, with all variables scaled to unit standard deviation. Table II shows that the accumulation tracks the realization in both halves of the morning. The uncertainty accumulated over 10:00–12:00 ET forecasts the announcement-window move beyond the 10:00 ET level, with a coefficient of 0.26 per standard deviation. The second window isolates the uncertainty accumulated in the two hours nearest the release, where a rising risk premium is most plausible. That increment still forecasts, with a coefficient of 0.15 per standard deviation, and the prediction for the press conference repeats both results (0.26 and 0.17), both significant at the 1% level. The pre-announcement accumulation is therefore not an increasing premium on an expectation formed early in the day. Instead, the market revises the expected impact of the announcement upward hour by hour, while the content of the decision remains undisclosed until the release.

5.1.5 Disappearance of the Pre-Announcement Drift

Lucca and Moench (2015) document that the S&P 500 gains about 49 basis points on average over the 24 hours before scheduled FOMC announcements, across their 1994–2011 sample. They document a negative co-movement between the pre-announcement returns and changes in the VIX. Hu et al. (2022) read the falling VIX as the resolution of the higher-order uncertainty of the impact, investors’ uncertainty about how large the announcement’s impact will be, the variance of the event uncertainty measured in this paper. However, the pre-announcement drift has *disappeared* in the later years.

Figure 5 plots the average cumulative ES futures return and the average cumulative VIX change over the two days around the announcement. The celebrated climb is absent. Where the drift in earlier samples is visible through the pre-announcement session, the 2016–2023

cumulative return is almost flat through both sessions and statistically insignificant at every minute before 14:00 ET. The same happens for the VIX: its pre-announcement decline disappears as well, with the average cumulative change statistically insignificant. Table III quantifies the disappearance window by window. The widest window, from the previous day’s open to the release, earns a cumulative return of 0.18% with a standard error of 0.12%, which is statistically insignificant. Over the standard 24-hour window of Lucca and Moench (2015), from 14:00 ET the day before to the release, the mean is 0.21% and significant at the 5% level. But it cannot be attributed to the pre-announcement trading session. The previous afternoon contributes 0.05%, the overnight gap from the preceding close to the announcement-day open earns 0.19%, and the announcement-day session, from the open to the release, only earns −0.03%. Only the overnight gap is significant, and the significant 24-hour return is therefore concentrated only in the overnight gap. The VIX shows the same pattern: its only significant decline, −0.85%, sits in the overnight gap, and its change is insignificant over every other window. Figures A.IV and A.V and Table A.II in the Appendix repeat the analysis with the VIX built on calendar rather than stochastic maturity, with the same conclusions.

The disappearance is consistent with the decline the literature documents in stages. Boguth, Grégoire, and Martineau (2019) find that over 2011–2017 the drift survives only ahead of the meetings followed by a press conference, and read the pattern as investor attention concentrating on the meetings expected to matter. Since 2019 every meeting carries a press conference, so the attention channel predicts a drift at each one, but none appears. Kurov, Wolfe, and Gilbert (2021) extend the sample and show the drift essentially disappears over 2016–2019, with or without a press conference.

These subsample comparisons bracket the disappearance within a multi-year window. Figure 6 implies there is a break in the drift at a single date, rather than a gradual decline over several years. Treating each announcement of the 2011/04–2023/12 sample as a candidate break date c , I take the 24 announcements around it, the 12 before and the 12 from c on, and estimate the shift in the average drift at c ,

$$r_i^{\text{pre}} = \alpha + \beta \mathbf{1}(i \geq c) + \varepsilon_i, \quad (32)$$

where r_i^{pre} is the pre-announcement drift of announcement i , the cumulative log return of the ES futures from 09:30 ET of the preceding trading day to the release. Among all

announcements, the largest shift with $\hat{\beta} = -0.76\%$, significant at the 1% level, happens on the March 2016 announcement, the first announcement after Cboe listed the Wednesday SPX expirations on February 23, 2016. The neighboring candidates from January to November 2016 also show the break with significant negative shifts, as their windows straddle the introduction; no other candidate in the sample produces a significant downward shift. The right column repeats the test on the announcements followed by a press conference, where the quarterly schedule before 2019 extends the regression window over roughly six years. The significant downward shifts again surround the introduction of the Wednesday expirations, spread from late 2015 to 2017 as the long windows imply.

The collapse at the break is abrupt. Figure 7 plots the average cumulative return path over the two-day window, for the nearest 4, 8, 16, and all announcements before and after the introduction of Wednesday-expired 0DTE options from 2012. Before the introduction, the pre-announcement drift is positive and statistically significant in all samples, from 0.27% over all announcements to 1.05% over the nearest four announcements. After the introduction, the average drift is between -0.13% and 0.20% , and statistically insignificant in all samples. Even across the nearest four announcements, the drift decreases by 1.01% to 0.04% and is statistically insignificant. The disappearance is not a gradual decline over several years, but a one-meeting collapse at the introduction of the Wednesday-expired 0DTE options. Figure 8 restricts the comparison to the announcements followed by a press conference, where Boguth, Grégoire, and Martineau (2019) locate the surviving drift. For these announcements, the average pre-announcement drift still falls from 0.47%–0.62% before the introduction to 0.14%–0.26% after it, and overall the drift of post-introduction announcements with a press conference is statistically insignificant.

Hu et al. (2022) imply an explanation for why the options listing can end the drift. In their two-period model, the pre-announcement drift is compensation for bearing the *uncertainty of event uncertainty*, which is resolved before the announcement. The same model predicts that the drift and the falling VIX disappear together once the event uncertainty is known in advance and no higher-order uncertainty is left to price. The Wednesday expirations cover the weekday of the FOMC announcements. The 0DTE options on the FOMC Wednesdays provide a direct measure of the event uncertainty. From then on, the event uncertainty is no longer a latent quantity investors must guess at, and the option market publishes it minute by minute. With the event uncertainty continuously observable and tradable, uncertainty

about it, the higher-order object priced in [Hu et al. \(2022\)](#), has little room to build, and the pre-announcement drift disappears with it.

5.1.6 Event Signal across Maturities

Comparing across maturities shows how the event signal is concentrated in the 0DTE options. Since FOMC announcements occur on Wednesdays, the announcement-day comparison uses the 0DTE, 2DTE, and 5DTE options available throughout the sample.⁵ Figure 9 plots the ATM implied volatility across maturities together with the spot volatility SV, for the announcement days and the days without FOMC announcements, where the non-FOMC column covers the full set from 0DTE to 5DTE. On announcement mornings the 0DTE ATM implied volatility climbs from 22% to 32%, while the 2DTE rises by less than one percentage point (17.1% to 17.8%) and the 5DTE by less than half a point. The accumulating event variance is, at any instant, a fixed quantity spread over the option’s remaining life, and two extra days of maturity already shrink its share enough to make the same event nearly invisible in the level. After the announcement, the ATM implied volatilities with longer maturities fall below spot volatility. The spot volatility SV jumps to 23% and stays elevated through the afternoon, yet the 2DTE and 5DTE implied volatilities sit 6% to 7% *below* it, while the 0DTE one remains above. Most of the remaining life of the longer options lies after the event, when the turbulence is expected to decay, so they price that calmer average rather than the current level. A longer maturity therefore weakens the event measure in two ways. Before delivery it shrinks the accumulating event variance to a small share of what it prices, and after delivery it averages in the post-event calm.

Days without FOMC announcements show none of this structure, and the absence is informative. The right column of Figure 9 shows the ATM implied volatilities with maturities from 0DTE to 5DTE lie within 0.8% of one another through 15:00 ET, about 3% *above* SV on the non-FOMC days. Each maturity traces a shallow U-shape, dipping into midday and flattening through the afternoon, with the deepest dip at the 0DTE. The alignment is the stochastic maturity of Section 3.1 at work. Option prices identify total variance, not volatility and maturity separately, so whatever the clock misstates about how variance accrues

⁵On the FOMC Wednesdays, the 1DTE and 3DTE options are excluded because the Tuesday and Thursday expirations are only available after April 2022. The non-FOMC panel pools all trading days and therefore shows every maturity from 0DTE to 5DTE.

reappears as implied-volatility structure. [Bondarenko \(2026\)](#) illustrates that under calendar time, which counts slow overnight and weekend hours at full weight, the USM term structure oscillates mechanically from one expiry to the next. He then shows the same alignment when using the stochastic clock: the mean SPX implied volatilities at the one-, three-, and five-day maturities stay within 1.3% of one another. The intraday shape depends on the clock as well. Figure [A.VI](#) in the Appendix repeats the non-FOMC panels on the same option prices, and calendar maturity makes the non-0DTE ATM implied volatility a mechanical decline through the day and business maturity makes it a mechanical rise, while the stochastic clock removes both trends and leaves the shallow U-shape except at the 5DTE. From 15:00 ET the 0DTE turns up, gaining 1.1% by 15:40 ET as the session approaches the closing auction, and the rise appears under all three clocks. The turn fades with maturity, a modest late recovery at the 1DTE and 2DTE, little at the 3DTE and 4DTE, and a slight drift lower at the 5DTE. The close, unlike the announcement, creates a much smaller wedge across maturities. A scheduled announcement enters each option’s remaining life only once, so its variance share dies off at rate $1/T$ at each timestamp across maturities, whereas the closing auction recurs every day: an option with N days of life contains N closes, so its share does not shrink with maturity. Only the 0DTE, for which today’s close is the single largest component of remaining risk, rises visibly more. The details are analyzed in [Section 6](#).

[Table IV](#) provides further support using predictive regressions. To treat the maturities symmetrically, I use the raw ATM implied volatility of each maturity, with no event extraction imposed, and predict the realized discontinuous volatility, JV of [equation \(19\)](#), to examine how much of the realized event volatility is embodied in the option at each maturity. I pool observations across announcement days and all sampling times s before 14:00 ET with intraday timestamp fixed effects, holding the forecast window fixed within each day, and estimate

$$\widetilde{JV}_{t,t+h} = \beta_I \widetilde{I}_s(0, T) + \beta_{SV} \widetilde{SV}_s + \widetilde{\varepsilon}_s, \quad (33)$$

where $[t, t+h)$ is the forecast window fixed at either the announcement ($[14:00, 14:15)$) or the press-conference ($[14:30, 15:00)$) window, and the tilde denotes the deviation from the timestamp- s mean, which absorbs one-way intraday timestamp fixed effects and removes the intercept. The explanatory power falls from 78% with 0DTE options to 67% with 2DTE and 58% with 5DTE for the announcement window, with a parallel decline (72%, 66%, 55%) for the press-conference window. The decline is a loss of event information. Spot volatility

alone explains 39% (32%) of the announcement (press-conference) window, so the margin the option adds over the persistence benchmark halves from 39% of R^2 at 0DTE to 20% at 5DTE. Each added day of maturity adds another day of background risk to the option price, so the event becomes a smaller share of what the instrument prices and the signal degrades. The maturity gradient implies the remaining disagreement with [Johannes, Kaeck, and Seeger \(2026\)](#), over the pre-announcement trend of the event uncertainty, is not a matter of methodology but of maturity. [Johannes, Kaeck, and Seeger \(2026\)](#) estimate one event volatility parameter for all maturities from the cross-section of options with up to 60 days to maturity, and find no average trend in it from two days before the announcement through 13:00 ET on the announcement day. The binding constraint is the common background state across maturities. Because the event’s share shrinks with maturity, the mix of the remaining risks differs across maturities. One common set of background parameters cannot fit every mix at once, and thus the variation it misses is attributed to the event parameter. However, at the 0DTE end of the maturity spectrum, where the event dominates the price, the event component unambiguously builds.

5.2 FOMC Minutes Release

The FOMC Minutes, released at 14:00 ET three weeks after each meeting, disclose the deliberations behind a decision the market already knows: a scheduled event with the same timing structure as the announcement but a fraction of its information content. [Figure 10](#) shows that the announcement-day pattern persists on the Minutes days, though at a smaller scale. The left column plots the intraday patterns of the measurements from [Section 4](#). Realized volatility RV drifts down from 16% to 10% before the release, then spikes to 27% at 14:00 ET, about one third of the peak on announcement days. The wedge between the 0DTE ATM implied volatility and SV widens from 2% to 8% before the release, and the event uncertainty $\hat{U}^A$ climbs from 0.25% to 0.34%, then collapses to 0.19% after the document becomes public. The peak event uncertainty is less than half its announcement-day level, so the option market prices the Minutes as the same kind of scheduled information delivery, carrying less of it.

The right column shows the results of the predictive regressions of [equation \(28\)](#) on the Minutes release days. The explanatory power rises to 70%–75% once the forecast window

contains or follows the release. The significance of the coefficients of $\widehat{U}_s^A$ is still concentrated in the windows containing the release, but only for sampling times s after roughly 11:00 ET, whereas on announcement days the predictive content is available from any point in the morning. When the event variance is a smaller share of what the option prices, the signal is noisier and decays faster. That a lower-information event reproduces the accumulation-and-release structure, scaled to its content, supports the reading that event uncertainty accumulates until delivery rather than resolving early.

6. Days without Scheduled Events

This section discusses the expected discontinuous activity on days without scheduled events, either because the day carries none or because the events have already passed. The sample also includes the intraday windows on FOMC announcement and Minutes days after the events have passed. The sample therefore consists of the 15:00–16:00 ET window on the FOMC announcement days and the 14:15–16:00 ET window on the Minutes release days, and all days that are neither FOMC announcement days nor Minutes release days, excluding the COVID period (March 2020) when the market was unusually volatile.

6.1 Intraday Patterns of the Discontinuity

Even without scheduled news, discontinuous risk does not spread evenly across the session. Table V reports ten-minute RV and JV on non-event days, alongside the post-event afternoons of the FOMC days, and documents a concentration of discontinuous activity at the close. Over 10:00–15:40 ET the realized JV averages 19%; it rises 5% to 24% in the pre-close window [15:40, 15:50) and then jumps to 36% in the closing window [15:50, 16:00), 1.9 times the previous average, and the RV increases from 48% to 63%. The same increase holds after the events pass on the FOMC days. The realized JV rises from 20% to 36% into the close on Minutes afternoons and from 23% to 38% on announcement afternoons.

The concentration at the close is not a surprise. At 15:50 ET the exchanges begin disseminating indicative closing-auction prices and order imbalances, market-on-close order deadlines pass, and the final ten minutes host a regularly anticipated burst of price discovery (Jegadeesh and Wu, 2022). As indexing and exchange-traded funds channel benchmark-

tracking flow to the closing price, the auction’s trading volume has increased in recent years (Jegadeesh and Wu, 2022; Bogousslavsky and Muravyev, 2023). Closing prices respond to the order imbalances of the final minutes, by amounts that depend on the auction’s design and its capacity to absorb late orders, which is genuinely uncertain until the end (Hu and Murphy, 2026). The close is therefore a focal event, also anticipated and time-stamped, but with a trigger that is microstructural rather than informational. It is a scheduled but not instantaneous event, and therefore the underlying discontinuous moves are not concentrated at a single known instant but spread across the final minutes, arriving at a finite rate. Existing studies examine the close from within: the volume it draws, the efficiency of price discovery, the design that absorbs its imbalances; whether the option market anticipates this recurring concentration *before* it arrives is still open.

For such a spread-out event, the discontinuous variance is of the same order as the diffusive variance, so the level of the ATM implied variance can no longer isolate a distinct discontinuous volatility. What it identifies instead is the expected magnitude of discontinuous moves, $\widehat{\mathcal{M}}^{N,J}$, through equation (10), and the ATM slope adds the net directional component $\widehat{\mathcal{J}}^N$ and its signed parts $\widehat{\mathcal{J}}^{N+}$ and $\widehat{\mathcal{J}}^{N-}$ through equation (14). The identified quantity changes, but the economic object does not. The magnitude $\widehat{\mathcal{M}}^{N,J}$ measures the expected aggregate size of large and discontinuous moves, which co-moves with the volatility of the discontinuous components. Similarly, $\widehat{\mathcal{J}}^N$ measures the expected net moves, which also reflects the directional asymmetry of the discontinuous components. Table VI summarizes the implied measures on the same days and horizons. The expected magnitude $\widehat{\mathcal{M}}^{N,J}$ is largest at the shortest maturity and where recent news leaves the most to digest. The 0DTE mean runs 0.99 on non-event days, 1.51 on post-Minutes-release afternoons, and 1.83 after the press conference. The net directional component $\widehat{\mathcal{J}}^N$ stays negative on average, -0.14 on non-event days, $-0.25(-0.51)$ on post-Minutes-release (post-press-conference) afternoons. Across maturities the expected magnitude falls with the option’s remaining life, to only 0.21 at 5DTE on non-event days, because a 0DTE option’s afternoon narrows toward the closing auction while a longer maturity averages that burst against days of quieter trading.

Figure 11 zooms in on the intraday patterns of JV and the 0DTE implied measures. In the top row, realized JV runs close to its session average until the final minutes and then steps up inside the closing window, peaking at the highest level of the day. The middle row shows the option side. The expected magnitude $\widehat{\mathcal{M}}^{N,J}$ accumulates toward the close, much

as $\widehat{\mathcal{U}}^A$ accumulates toward 14:00 ET on FOMC event days, from 0.29 at 10:00 ET to 1.71 by 15:00 ET, then doubling to 3.05 at 15:40 ET. The post-event $\widehat{\mathcal{M}}^{N,J}$ rebuilds in the same way, from higher bases. No matter whether the day contained a scheduled event or nothing at all, the option market anticipates the close as a burst of discontinuous activity. In the bottom row both signed components widen into the close, from 0.07 to 1.36 for $\widehat{\mathcal{J}}^{N+}$ and from -0.22 to -1.69 for $\widehat{\mathcal{J}}^{N-}$ on non-event days, while the net component never leaves the band -0.14 to -0.33 . The market thus expects the close as a growing two-sided move with the usual left tilt, not as a directional bet.

6.2 Evidence from Predictive Regressions

The intraday accumulation is suggestive, but it does not establish whether the option market specifically anticipates the closing auction. I apply a predictive regression design similar to that of Section 5.1 and examine whether the expected discontinuous components specifically forecast the closing auction. The regression is pooled across sampling times s before 15:40 ET with intraday timestamp fixed effects,

$$\widetilde{\text{JV}}_{t,t+h} = \beta_{\mathcal{M}} \widetilde{\widehat{\mathcal{M}}_s^{N,J}} + \beta_{SV} \widetilde{SV}_s + \widetilde{\varepsilon}_s, \quad (34)$$

with the forecast window $[t, t+h)$ set to the pre-close window 15:40–15:50 ET or the closing window 15:50–16:00 ET. Table VII reports the results, across maturities from 0 to 5 DTE and for two samples, all non-event days and the normal days whose daily realized variance falls in the interquartile range of the sample. The predictive content is aimed at the auction, not at the end of the day in general. Every sampling time lies closer to the pre-close window than to the closing window, yet the nearer window is the harder one to forecast. At 0DTE the univariate R^2 of $\widehat{\mathcal{M}}_s^{N,J}$ is 8% for [15:40, 15:50) against 12% for [15:50, 16:00), and the bivariate regression reaches 29% against 43%. The gap widens with maturity. From 2DTE on, the pre-close coefficient of $\widehat{\mathcal{M}}_s^{N,J}$ is statistically insignificant, while the closing-window coefficients stay at 0.20–0.25 per standard deviation and significant at the 1% level. By 5DTE the pre-close content is essentially gone, a univariate R^2 of 0.4%, while the closing window retains 4%. Truncation at zero in (19) does not manufacture the gap between the two windows. In the pre-close window the ratio of zero JV is 7.4% against 2.7% for the closing window, so in both cases the regressions fit variation among the strictly positive

realizations that make up over 90% of the sample, not a mass at zero. What the option market anticipates is therefore not a generally turbulent final stretch but the discontinuous repricing time-stamped at the auction, even though a window ten minutes nearer sits in between.

The levels of predictive performance sit far below the event-day levels, as they should. Without an event at a known instant, the prospective discontinuity is a small share of what even a 0DTE option prices, and the closing burst competes with a full day of background volatility. The maturity gradient, in contrast, is far flatter than on announcement days. There, the marginal predictive content the option adds over the persistence of SV halves from 0DTE to 5DTE, because the announcement enters each option’s remaining life exactly once and its variance share dies off as the maturity lengthens. However, the closing auction recurs, one burst for every day of remaining life, so lengthening the maturity does not push the event out of the instrument. The margin over the SV is 6.4% at 0DTE and still 6.7% at 5DTE, and the falling univariate share traces the averaging of the close over the option’s longer life, which also lowers its mean in Table VI, rather than a vanishing signal. Furthermore, the panel of normal days shows the prediction is not repackaged persistence. The interquartile filter strips out most of the calm and turbulent variation across days, cutting the univariate share of SV from 36% to 10% for the closing window, yet the marginal predictive content of $\widehat{\mathcal{M}}^{N,J}$ remains 8.1% at 0DTE.

For all its predictive content, the level regression does not yet say what drives the intraday climb of $\widehat{\mathcal{M}}^{N,J}$. It could be a rising risk premium, a mechanical consequence of the approach of expiry as maturity in the denominator shrinks, or a combination of both. The increment decomposition of equation (31) transfers directly, with the expected discontinuity magnitude in place of the event uncertainty,

$$\text{JV}_{t,t+h} = \alpha + \beta_{\Delta\mathcal{M}} \Delta\widehat{\mathcal{M}}_s^{N,J} + \beta_{\mathcal{M}} \widehat{\mathcal{M}}_{\text{ref}}^{N,J} + \beta_{SV} \text{SV}_s + \varepsilon_{t,t+h}, \quad (35)$$

where $\Delta\widehat{\mathcal{M}}_s^{N,J} = \widehat{\mathcal{M}}_s^{N,J} - \widehat{\mathcal{M}}_{\text{ref}}^{N,J}$ is the expected magnitude accumulated between the reference time and the sampling time s , and the forecast window $[t, t+h)$ is fixed at the closing window [15:50, 16:00) ET. Three accumulation windows split the session, 10:00–12:00, 12:00–14:00, and 14:00–15:40 ET, so the reference level absorbs everything priced before the window opens and the increment carries only what accumulates within it. Table VIII shows that

every increment forecasts the closing burst. At 0DTE the coefficient rises from 0.23 in the morning window to 0.30 in the window nearest the auction, the increments estimated from 0DTE to 2DTE options are all significant at the 1% level, and the 3DTE to 5DTE options in Appendix Table A.III behave alike. The normal-day panel sharpens the result: filtering out the calm and turbulent day types strengthens the coefficients of increments to 0.38–0.40 per standard deviation, the largest regressors after the spot volatility control. The expected magnitude of the closing repricing is revised upward through the day, and the revisions track the realization, the same calibrated accumulation that prices the scheduled events.

As the pooled and window designs compress the trading day into a few estimates per maturity, Figure 12 restores the clock. I re-estimate the same bivariate design at every sampling minute s from 10:00 to 15:40 ET on non-event days, using 0DTE options and the same forecast windows, each sampling minute producing a separate time-series regression across days,

$$JV_{t,t+h} = \alpha + \beta_{\mathcal{M}} \widehat{\mathcal{M}}_s^{N,J} + \beta_{SV} SV_s + \varepsilon_{t,t+h}, \quad (36)$$

$$r_{t,t+h} = \alpha + \beta_{\mathcal{J}^N} \widehat{\mathcal{J}}_s^N + \beta_{SV} SV_s + \varepsilon_{t,t+h}. \quad (37)$$

The closing auction is forecastable from anywhere in the day. The R^2 is already 42% at $s = 10 : 00$ ET and stays between 41% and 46% through the session, and the slope on $\widehat{\mathcal{M}}^{N,J}$ is positive and significant throughout. Predicting the pre-close window shows the same shape at lower levels, with R^2 of 26%–33%. The predictive content of the expected discontinuous magnitude is therefore a matter of the closing auction, not of the final hour in general, and the option market anticipates it from the open. Direction is again unpredictable. For the closing window the R^2 never exceeds 0.5%, and the coefficient on $\widehat{\mathcal{J}}^N$ is statistically insignificant at almost every sampling time. The option market thus anticipates the closing auction in the same way as the scheduled events. The uncertainty about the concentration of price discovery accumulates through the day, but the market learns nothing about its direction until it is delivered.

6.3 Post-Event Afternoons

Post-event afternoons tie the two scenarios together on the same calendar days. Table IX estimates the pooled regression on the horizons after the release, from 14:15 ET on Minutes

days and from 15:00 ET on announcement days, where the scheduled event no longer sits inside the option’s remaining life. On Minutes afternoons the expected discontinuous magnitude $\widehat{\mathcal{M}}^{N,J}$ is still a live predictor, with univariate R^2 of 23% and bivariate R^2 of 60% at 0DTE. Although most of the bivariate performance is shared with SV, which the residual from the release has elevated, the magnitude’s own increment is still 18.1%. On announcement afternoons the decaying turbulence of the press conference sits close to the auction and dominates the 0DTE option’s remaining life, leaving little predictive content for the closing window. Overall, a trading day is not one scenario or the other. It migrates from the scheduled-event identification to the non-event identification the moment the event happens, and the option market’s pricing switches with it, accumulating toward whichever concentration of price discovery comes next.

7. Conclusion

When the market knows that information is coming but not its content, uncertainty about the impact builds until the information is released rather than resolves in advance. I document this by reading the expected size and tilt of the prospective FOMC move from one-minute 0DTE option prices. The expected event uncertainty rises 30% through the morning, peaks as the statement crosses the wire, collapses, rebuilds ahead of the press conference, and collapses again once the chair begins speaking. The statement rather than the press conference resolves the majority of the priced event variance. The buildup of the event uncertainty carries genuine forecast content rather than only a risk premium. Sampled at any point in the morning, it explains 78%–85% of the cross-day variation in realized volatility, but only in forecast windows that contain the statement release or the press conference. The implied distribution keeps a stable left tilt, a skewness near -0.5 , that neither deepens nor weakens as the release approaches, so option prices carry no advance directional view of the decision. Moreover, the pre-FOMC announcement drift documented in the literature disappears in recent years, and the VIX decline is also gone. The disappearance of the drift dates to the first announcement after Cboe introduced the Wednesday-expiring 0DTE options, which is consistent with the model of [Hu et al. \(2022\)](#). Once the event uncertainty is observable and tradable, there is no longer any higher-order uncertainty to compensate, and the drift and VIX decline vanish together. Furthermore, the FOMC Minutes, a disclosure with identical

timing but a fraction of the information, reproduce the accumulation and release path at less than half the scale, with a noisier and later-arriving forecast signal.

Days without scheduled news carry no event at a known time, yet the same accumulation appears around the closing auction, an anticipated concentration of price discovery whose trigger is microstructural rather than informational. The identified quantity changes to the expected magnitude and net direction of discontinuous moves, but the economic content does not: the expected magnitude builds through the session toward the close and forecasts the closing-window realized discontinuity from any sampling point in the day, again with no directional content. What the option market prices in this accumulating fashion is therefore any regularly anticipated concentration of price discovery, whether the trigger is a macroeconomic announcement or the microstructure of the close.

REFERENCES

- Adams, Greg, Chukwuma Dim, Bjorn Eraker, Jean-Sebastien Fontaine, Chayawat Ornthanalai, and Grigory Vilkov, 2025, Do S&P500 Options Increase Market Volatility? Evidence from 0DTEs, *Evidence from 0DTEs (October 17, 2025)* .
- Ai, Hengjie, and Ravi Bansal, 2018, Risk Preferences and the Macroeconomic Announcement Premium, *Econometrica* 86, 1383–1430.
- Ai, Hengjie, Ravi Bansal, and Leyla Jianyu Han, 2021, Information Acquisition and the Pre-announcement Drift, *Available at SSRN 3964349* .
- Ai, Hengjie, Leyla Jianyu Han, and Lai Xu, 2026, Information-driven Volatility, *Available at SSRN 3961096* .
- Almeida, Caio, Gustavo Freire, and Rodrigo Hizmeri, 2024, 0DTE Asset Pricing, *Available at SSRN 4701401*.
- Anderson, Torben G, Tim Bollerslev, Francis X Diebold, and Clara Vega, 2003, Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review* 93, 38–62.
- Backus, David, Silverio Foresi, Kai Li, and Liuren Wu, 1997, Accounting for Biases in Black-Scholes, *CRIF Working Paper, November 1997* .

- Bandi, Federico M, Nicola Fusari, and Roberto Renò, 2024, 0dte option pricing, *Available at SSRN 4503344*.
- Beckmeyer, Heiner, Nicole Branger, and Leander Gayda, 2023, Retail Traders Love 0DTE Options... But Should They?, *Available at SSRN 4404704* .
- Beckmeyer, Heiner, Nicole Branger, and Thomas Grünthaler, 2021, Fed Tails: FOMC Announcements and Stock Market Uncertainty, Working paper, University of Münster.
- Bernile, Gennaro, Jianfeng Hu, and Yuehua Tang, 2016, Can Information Be Locked Up? Informed Trading Ahead of Macro-News Announcements, *Journal of Financial Economics* 121, 496–520.
- Bougousslavsky, Vincent, and Dmitriy Muravyev, 2023, Who Trades at the Close? Implications for Price Discovery and Liquidity, *Journal of Financial Markets* 66, 100852.
- Boguth, Oliver, Vincent Grégoire, and Charles Martineau, 2019, Shaping Expectations and Coordinating Attention: The Unintended Consequences of FOMC Press Conferences, *Journal of Financial and Quantitative Analysis* 54, 2327–2353.
- Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *The Journal of Finance* 66, 2165–2211.
- Bollerslev, Tim, Viktor Todorov, and Sophia Zengzhi Li, 2012, Jump Tails, Extreme Dependencies, and the Distribution of Stock Returns, *Journal of Econometrics* forthcoming, unknown.
- Bondarenko, Oleg, 2026, The Clock Matters: Time Measurement and the Ultra-Short Volatility Term Structure, *Available at SSRN 6705119* .
- Brogaard, Jonathan, Jaehee Han, and Peter Y Won, 2023, Does 0DTE Options Trading Increase Volatility?, *Available at SSRN 4426358* .
- Chong, Carsten H., and Viktor Todorov, 2024, Volatility of Volatility and Leverage Effect from Options, *Journal of Econometrics* 240, 105669.
- Chong, Carsten H, and Viktor Todorov, 2025, Do Equity and Options Markets Agree about Volatility?, *Journal of Finance forthcoming* .
- Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen, 2019, Stock Returns over the FOMC Cycle, *The Journal of Finance* 74, 2201–2248.
- Cocoma, Paula, 2025, Disagreement and Scheduled Announcements: Explaining the Pre-Announcement Drift, Working paper, Frankfurt School of Finance and Management.

- Cushing, David, and Ananth Madhavan, 2000, Stock Returns and Trading at the Close, *Journal of Financial Markets* 3, 337–365.
- Dorion, Christian, Piotr Orłowski, and Yuhan Song, 2026, The Factor Structure of 0DTE Option Returns, *Working paper* .
- Dubinsky, Andrew, Michael Johannes, Andreas Kaeck, and Norman J Seeger, 2019, Option Pricing of Earnings Announcement Risks, *The Review of Financial Studies* 32, 646–687.
- Ederington, Louis H, and Jae Ha Lee, 1993, How Markets Process Information: News Releases and Volatility, *The Journal of Finance* 48, 1161–1191.
- Fleming, Michael J, and Eli M Remolona, 1999, Price Formation and Liquidity in the US Treasury Market: The Response to Public Information, *The journal of Finance* 54, 1901–1915.
- François, Pascal, and Lars Stentoft, 2021, Smile-Implied Hedging with Volatility Risk, *Journal of Futures Markets* 41, 1220–1240.
- Hall, Peter, 2013, *The Bootstrap and Edgeworth Expansion* (Springer Science & Business Media).
- Hu, Edwin, and Dermot Murphy, 2026, Vestigial Tails? Floor Brokers at the Close in Modern Electronic Markets, *Management Science* 72, 3974–3996.
- Hu, Grace Xing, Jun Pan, Jiang Wang, and Haoxiang Zhu, 2022, Premium for Heightened Uncertainty: Explaining Pre-announcement Market Returns, *Journal of Financial Economics* 145, 909–936.
- Jegadeesh, Narasimhan, and Yanbin Wu, 2022, Closing Auctions: Nasdaq versus NYSE, *Journal of Financial Economics* 143, 1120–1139.
- Johannes, Michael Slater, Andreas Kaeck, and Norman Seeger, 2026, FOMC Announcement Event Risk, *Journal of Finance*, *forthcoming* .
- Kurov, Alexander, Marketa Halova Wolfe, and Thomas Gilbert, 2021, The disappearing pre-FOMC announcement drift, *Finance Research Letters* 40, 101781.
- Laarits, Toomas, 2025, Pre-Announcement Risk, Working paper, New York University Stern School of Business; first draft November 15, 2018.
- Lucca, David O, and Emanuel Moench, 2015, The Pre-FOMC Announcement Drift, *The Journal of finance* 70, 329–371.
- Medvedev, Alexey, and Olivier Scaillet, 2007, Approximation and Calibration of Short-Term

- Implied Volatilities Under Jump-Diffusion Stochastic Volatility, *Review of Financial Studies* 20 (2), 427–459.
- Savor, Pavel, and Mungo Wilson, 2013, How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements, *Journal of Financial and Quantitative Analysis* 48, 343–375.
- Todorov, Viktor, 2019, Nonparametric spot volatility from options, *The Annals of Applied Probability* 29, 3590–3636.
- Vasquez, Aurelio, Diego Amaya, Neil D Pearson, and Pedro Angel Garcia-Ares, 2025, 0DTE Index Options and Market Volatility: How Large is Their Impact?, *Available at SSRN 5113405* .
- Wachter, Jessica A, and Yicheng Zhu, 2022, A Model of Two Days: Discrete News and Asset Prices, *The Review of Financial Studies* 35, 2246–2307.
- Yan, Shu, 2011, Jump Risk, Stock Returns, and Slope of Implied Volatility Smile, *Journal of Financial Economics* 99, 216–233.

Figure 1. Intraday Patterns of Realized and Implied Volatility on FOMC Announcement Days

This figure plots the intraday patterns of one-minute realized and implied ATM volatility from 10:00 to 15:00 ET on FOMC announcement days. Each column shows the subsample of different policy decisions: all FOMC announcement days, days with a rate change (Cut or Hike), and days with no change (Hold). The top panel plots the intraday patterns of realized volatility, RV, measured by the square root of realized variance computed from five-second ES futures returns. The middle panel plots the intraday patterns of the realized continuous volatility, CV, measured by the truncated realized variation that isolates the diffusive part, and the realized jump volatility, JV, measured by the square root of the difference between realized variance and continuous variance. The bottom panel plots the intraday patterns of 0DTE ATM implied volatility and the spot volatility, SV. The estimation of these variables is described in Section 4. All variables are annualized and expressed in percent. Vertical dashed lines mark 14:00 and 14:30 ET, the scheduled announcement and press conference times, respectively. The sample period is from 2016/09 to 2023/12.




Figure 2. Intraday Patterns of Expected Event Uncertainty and Skewness on FOMC Announcement Days

This figure plots the intraday patterns of one-minute expected event uncertainty and skewness from 0DTE options between 10:00 and 15:00 ET on FOMC announcement days. Each column shows the subsample of different policy decisions: all FOMC announcement days, days with a rate change (Cut or Hike), and days with no change (Hold). The top and bottom panels plot the intraday patterns of the expected event volatility, $\hat{U}^A$, and the expected event skewness, $\hat{S}^A$, respectively, both estimated from 0DTE options, described in Section 4. Vertical dashed lines mark 14:00 and 14:30 ET, the scheduled announcement and press conference times, respectively. The lines with lighter colors are the raw minute-level values, and the lines with darker colors are ten-minute centered rolling averages that reset at 14:00 and 14:30 ET. The sample period is from 2016/09 to 2023/12.




Figure 3. Predicting Realized Volatility and Returns on FOMC Announcement Days

This figure plots the R^2 and p -values of the predictive regressions (28) and (29) of realized volatility and log returns on one-minute expected event volatility and skewness, controlling for the spot volatility, respectively, on FOMC announcement days. The $\hat{U}_s^A$ and $\hat{S}_s^A$ are the expected event volatility and skewness estimated from 0DTE options, and SV is the spot volatility. The estimation of these variables is described in Section 4. The predictive regression is applied for every pair of the one-minute frequency time s (10:00–13:00 ET) and five-minute frequency forecast end time $t + 5\text{min}$ (13:05–15:00 ET). The left column plots the adjusted R^2 (bars) evaluated at $s = 13:00$ ET with the gray error bars showing the minimum and maximum adjusted R^2 across all s , and the x-axis is the forecast end time $t + 5\text{min}$. The middle column plots the heatmaps of the p -values of the coefficients on $\hat{U}_s^A$ and $\hat{S}_s^A$ on the top and bottom panel, respectively. Both heatmaps have the time s of the regressors on the y-axis and the corresponding forecast end time $t + 5\text{min}$ on the x-axis, with yellow cells denoting smaller p -values and blue cells denoting larger p -values, on the $[0, 0.1]$ scale. The right column plots the heatmaps of the p -values of the coefficients on the spot volatility control, SV, with the same axis and color scheme as the middle column. Each variable is scaled to unit standard deviation. The sample period is from 2016/09 to 2023/12.




Figure 4. Predicting Realized Volatility Using Alternative Event Uncertainty on FOMC Announcement Days

This figure plots intraday patterns of the alternative event volatility, and the R^2 and p -values of the predictive regression (28) of realized volatility on the one-minute alternative expected event uncertainty, controlling for the spot volatility, on FOMC announcement days. The $\hat{U}^{A,alt}$ is the alternative expected event uncertainty estimated from 0DTE and 2DTE options, as in equation (30), and SV is the spot volatility, described in Section 4.2. The predictive regression is applied for every pair of the one-minute frequency time s (10:00–13:00 ET) and five-minute frequency forecast end time $t + 5\text{min}$ (13:05–15:00 ET). The left column plots the intraday patterns of the alternative expected event volatility, $\hat{U}^{A,alt}$, shown as lighter lines for the raw minute-level values and darker lines for the ten-minute rolling mean, with the vertical dashed lines marking 14:00 and 14:30 ET, the scheduled announcement and press conference times, respectively. The middle column plots the adjusted R^2 (bars) evaluated at $s = 13 : 00$ ET with the gray error bars showing the minimum and maximum adjusted R^2 across all s , and the x-axis is the forecast end time $t + 5\text{min}$. The right column plots the heatmap of the p -values of the coefficient on $\hat{U}^{A,alt}$. The heatmap has the time s of the regressors on the y-axis and the corresponding forecast end time $t + 5\text{min}$ on the x-axis, with yellow cells denoting smaller p -values and blue cells denoting larger p -values, on the $[0, 0.1]$ scale. Each variable is scaled to unit standard deviation. The sample period is from 2016/09 to 2023/12.

45




Figure 5. Pre-FOMC Announcement Drift and VIX Change

The figure plots the cumulative returns of the ES futures and the cumulative change of the VIX over the two-day window around FOMC announcements. The top row shows the average cumulative log return of the ES futures, accumulated from 09:30 ET of the trading day preceding the announcement. The bottom row shows the average cumulative log change of the 30-day VIX from the same starting point, where the VIX is computed from SPX options at the one-minute frequency following the standard Cboe methodology, with maturities measured by the stochastic maturity of Section 3.1. The columns split the announcements into two samples: all announcements and announcements after COVID (from 2020/04). Shaded areas are the 95% confidence bands across announcements, and the vertical dashed line marks the 14:00 ET announcement time. The sample period is from 2016/09 to 2023/12.




Figure 6. Pre-FOMC Announcement Drift Break Test

The figure tests for a break in the pre-FOMC announcement drift at every candidate meeting. For each announcement, the drift is the cumulative log return of the ES futures, from 09:30 ET of the trading day preceding the announcement to the announcement time (14:00 ET from 2013/03; 12:30 or 14:15 ET before). For each candidate meeting, the local sample contains 24 announcements, the 12 announcements before the candidate and the 12 announcements from the candidate on, and the drift is regressed on an indicator for the latter group. The top row plots the estimated shift $\hat{\beta}$ in the average drift, and the bottom row plots its p -value based on Newey–West standard errors. The left column uses all announcements, and the right column uses the announcements followed by a press conference (PC). Filled markers denote p -values at or below 10%, and the horizontal dashed lines mark zero in the top row and the 10% and 5% significance levels in the bottom row. The vertical dashed line marks the introduction of the Wednesday-expiring 0DTE options on February 23, 2016. The sample period is from 2011/04 to 2023/12.




Figure 7. Pre-FOMC Announcement Drift around Wednesday Expiration Introduction

The figure plots the average cumulative log return of the ES futures over the two-day window around FOMC announcements, comparing the announcements before and after the introduction of the Wednesday-expiring 0DTE options on February 23, 2016. Each path is accumulated from 09:30 ET of the trading day preceding the announcement. The left column averages the announcements before the introduction, and the right column averages the announcements after it. The rows widen the comparison window: the nearest 4, 8, and 16 announcements on each side of the introduction, and all announcements from 2011/04. Shaded areas are the 95% confidence bands across announcements, and the vertical dashed lines mark the announcement time (14:00 ET from 2013/03; 12:30 or 14:15 ET before). The sample period is from 2011/04 to 2023/12.




Figure 8. Pre-FOMC Announcement Drift Followed by Press Conference around Wednesday Expiration Introduction

The figure plots the average cumulative log return of the ES futures over the two-day window around FOMC announcements, comparing the announcements before and after the introduction of the Wednesday-expiring options on February 23, 2016. The announcements are restricted to the announcements followed by a press conference. Each path is accumulated from 09:30 ET of the trading day preceding the announcement. The left column averages the announcements before the introduction, and the right column averages the announcements after it. The rows widen the comparison window: the nearest 4, 8, and 16 announcements on each side of the introduction, and all announcements from 2011/04. Shaded areas are the 95% confidence bands across announcements, and the vertical dashed lines mark the announcement time (14:00 ET from 2013/03; 12:30 before). The sample period is from 2011/04 to 2023/12.




Figure 9. Intraday Patterns of ATM Implied Volatility across Maturities

The figure plots the intraday patterns of ATM implied volatility across maturities and spot volatility on FOMC announcement days and non-FOMC days. The left column shows FOMC announcement days from 10:00 to 15:00 ET for 0DTE, 2DTE, and 5DTE options, with vertical dashed lines marking 14:00 and 14:30 ET, the scheduled release and press conference times, respectively. The right column shows non-FOMC days, the days without FOMC announcements, from 10:00 to 15:40 ET for 0DTE to 5DTE options. Within each column, the top row plots the ATM implied volatility of all maturities and the spot volatility, SV, described in Section 4.2. The bottom row zooms in on the non-0DTE maturities to show the intraday patterns clearly. The sample period is from 2016/09 to 2023/12.




Figure 10. Intraday Patterns of 0DTE Implied Discontinuity Factors and the Predictive Power on FOMC Minutes Release Days

This figure plots the intraday patterns of 0DTE option-implied volatility measures and the predictive regression (28) on FOMC Minutes release days. The left column shows intraday time-series patterns over 10:00–14:15 ET and the right column shows predictive regression results. In the left column, the top panel plots realized volatility RV, the square root of one-minute realized variance built from five-second ES futures log returns; the middle panel plots the 0DTE ATM implied volatility $\hat{I}_t(0, T)$ and the spot volatility, SV; and the bottom panel plots the implied event uncertainty, $\hat{U}^A$, shown as raw minute-level values (faded) and as a rolling mean (window of ten minutes, computed separately before and after the 14:00 ET break). The estimation of these variables is described in Section 4. For the right column, the predictive regression is applied for every pair of the one-minute frequency time s (10:00–13:00 ET) and five-minute frequency forecast end time $t + 5\text{min}$ (13:05–14:15 ET). The top row plots the adjusted R^2 (bars) evaluated at $s = 13 : 00$ ET with the gray error bars showing the minimum and maximum adjusted R^2 across all s in 10:00–13:00 ET, and the x-axis is the forecast end time $t + 5\text{min}$. The middle and bottom panels plot the heatmaps of the p -values of the coefficients on $\hat{U}_s^A$ and SV, respectively. Both heatmaps have the time s of the regressors on the y-axis and the corresponding forecast end time $t + 5\text{min}$ on the x-axis, with yellow cells denoting smaller p -values and blue cells denoting larger p -values, on the $[0, 0.1]$ scale. Each variable is scaled to unit standard deviation. The sample period is from 2016/09 to 2023/12.




Figure 11. Intraday Patterns of 0DTE Implied Discontinuity Components on Non-Event, FOMC Minutes, and FOMC Announcement Days

This figure plots the intraday patterns of one-minute realized and 0DTE option-implied discontinuity factors on non-event days and during the post-event horizons on FOMC Minutes release and FOMC announcement days. Non-event days are the days excluding FOMC announcement and Minutes days and the COVID period (2020/03), with the intraday period from 10:00 ET to 15:40 ET. The post-event windows are from 14:15 ET to 15:40 ET for Minutes days and from 15:00 ET to 15:40 ET for FOMC announcement days. The top row plots realized jump volatility JV, from equation (19), with the red horizontal lines marking the average level over the session before 15:50 ET and the closing auction window, between 15:50 ET and 16:00 ET. The middle row plots the expected discontinuity magnitude $\widehat{\mathcal{M}}^{N,J}$, and the bottom row plots the expected net directional discontinuity $\widehat{\mathcal{J}}^N$, with upside and downside components $\widehat{\mathcal{J}}^{N+}$ and $\widehat{\mathcal{J}}^{N-}$. The estimation of these variables is described in Section 4. The sample runs from 2016/09 to 2023/12.



Figure 12. Intraday Patterns of Predictive Regressions for Discontinuity Volatility and Returns on Non-Event Days

This figure plots the R^2 and p -values of the predictive regressions (36) and (37) of realized discontinuity volatility and log returns on one-minute expected discontinuity magnitude and net direction from 0DTE options, controlling for the spot volatility, respectively, on non-event days. Non-event days are the days excluding FOMC announcement and Minutes days and the COVID period (2020/03), with the intraday period from 10:00 ET to 15:40 ET. The $\widehat{\mathcal{M}}^{N,J}$ and $\widehat{\mathcal{J}}^N$ are the expected discontinuity magnitude and net direction, and SV is the spot volatility. The estimation of these variables is described in Section 4. The predictive regression is applied for every sampling timestamp s from 10:00 ET to 15:40 ET and the forecast window is fixed at [15:40, 15:50) ET and [15:50, 16:00) ET for realized discontinuity volatility and log returns, respectively. The top row plots the adjusted R^2 at each timestamp s , the middle row plots the estimated coefficient on $\widehat{\mathcal{M}}^{N,J}$ and $\widehat{\mathcal{J}}^N$ for the left and right column, respectively, and the bottom row plots the corresponding p -values. Each variable is scaled to unit standard deviation. The sample period is from 2016/09 to 2023/12.




Table I
Summary Statistics of Expected Event Uncertainty Increments on FOMC Announcement Days

This table reports the mean and standard deviation of the pre-announcement increments in the expected event uncertainty and the average of its intraday trend on the FOMC announcement days. The increment of the expected event uncertainty, $\Delta\hat{\mathcal{U}}_s^A = \hat{\mathcal{U}}_s^A - \hat{\mathcal{U}}_{\text{ref}}^A$, is the change in the expected event volatility from the reference time to the sampling time s , from 0DTE options, described in Section 4 and expressed in percent. The columns $\Delta\hat{\mathcal{U}}^A(10:00\text{--}12:00)$ and $\Delta\hat{\mathcal{U}}^A(12:00\text{--}13:59)$ set the reference time at 10:00 ET and 12:00 ET, and the sampling time at $s = 12:00$ ET and 13:59 ET, respectively. The columns Slope $\times 60$ report the intraday trend of the expected event uncertainty. Each announcement day yields one slope coefficient from an OLS regression of the $\hat{\mathcal{U}}^A$ on a linear time trend over 10:00–13:59 ET at one-minute frequency, and the reported coefficient is the average of the daily slopes across announcement days, scaled to percentage per hour (multiplied by 60). The sample is grouped by the policy decision: all FOMC announcement days (All), days with a rate change (Cut/Hike), and days with no change (Hold). The standard errors and significance tests are based on Newey–West standard errors, with the announcement day as the unit of observation. The column N reports the number of announcement days in each group. The sample period is from 2016/09 to 2023/12.

Decision	$\Delta\hat{\mathcal{U}}^A(10:00\text{--}12:00)$		$\Delta\hat{\mathcal{U}}^A(12:00\text{--}13:59)$		Slope $\times 60$		N
	Mean	SD	Mean	SD	Coef.	SE	
All	0.10***	0.18	0.03**	0.09	0.038***	(0.007)	56
Cut/Hike	0.12**	0.20	0.01	0.07	0.050***	(0.013)	22
Hold	0.09***	0.17	0.04*	0.10	0.031***	(0.005)	34

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Table II
Predicting Event Volatility Using Uncertainty Increments on FOMC
Announcement Days

This table reports the coefficients and R^2 of the predictive regression (31) of the realized event volatility on the increment of event uncertainty, the level of the expected uncertainty at a reference time, and spot volatility on the FOMC announcement days. The realized event volatility, $JV_{t,t+h}$, is measured by the difference between the realized volatility and the continuous volatility, as in equation (19), over the window $[t, t+h)$. The regressors include the increment of the expected event uncertainty, $\Delta\hat{\mathcal{U}}_s^A = \hat{\mathcal{U}}_s^A - \hat{\mathcal{U}}_{\text{ref}}^A$, which is the change in the expected event volatility from the reference time to the sampling time s , from 0DTE options. The columns $\Delta(10:00\text{--}12:00)$ and $\Delta(12:00\text{--}13:59)$ set the reference time at 10:00 ET and 12:00 ET, and the sampling time at $s = 12:00$ ET and 13:59 ET, respectively. The controls are the level of the expected event uncertainty at the reference time, $\hat{\mathcal{U}}_{\text{ref}}^A$, from 0DTE options, and the spot volatility, SV_s , described in Section 4. Each column is a separate time-series regression across announcement days, holding the forecast window $[t, t+h)$ fixed, with all variables scaled to unit standard deviation. The predictive regression is applied for two forecast windows: [14:00, 14:15) ET, the 15 minutes after the announcement, and [14:30, 15:00) ET, the half hour after the press conference starts at 14:30 ET. The parentheses in the table report the Newey–West standard errors. The foot of each panel reports the R^2 (in %) and the number of announcement days N . The sample period is from 2016/09 to 2023/12.

JV(14:00–14:15)		
Factor	$\Delta(10:00\text{--}12:00)$	$\Delta(12:00\text{--}13:59)$
α	0.11 (0.16)	0.02 (0.10)
$\Delta\hat{\mathcal{U}}_s^A$	0.26*** (0.07)	0.15*** (0.05)
$\hat{\mathcal{U}}_{\text{ref}}^A$	0.73*** (0.09)	0.71*** (0.06)
SV_s	0.15 (0.09)	0.28*** (0.06)
R^2	76.91	80.55
N	56	56
JV(14:30–15:00)		
Factor	$\Delta(10:00\text{--}12:00)$	$\Delta(12:00\text{--}13:59)$
α	−0.11 (0.15)	−0.22* (0.12)
$\Delta\hat{\mathcal{U}}_s^A$	0.26*** (0.10)	0.17*** (0.05)
$\hat{\mathcal{U}}_{\text{ref}}^A$	0.75*** (0.10)	0.78*** (0.08)
SV_s	0.06 (0.06)	0.15* (0.08)
R^2	71.62	74.72
N	56	56

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Table III
Summary Statistics of the Pre-Announcement Drift and VIX Change

This table reports the mean and standard deviation of the pre-announcement cumulative log returns of the ES futures and the cumulative change of the VIX on FOMC announcement days. The $\Delta r_{t,t+h}$ is the cumulative log return of the ES futures from the start of the window t to its end time $t+h$. The $\Delta \log \text{VIX}_{t,t+h}$ is the cumulative log change of the 30-day VIX over the same window, in percent, where the VIX is computed from SPX options following the standard Cboe methodology, with maturities measured by the stochastic maturity of Section 3.1. The windows Prev. 09:30–14:00 and Prev. 14:00–14:00 start at 09:30 ET and 14:00 ET of the trading day preceding the announcement, end at the 14:00 ET release, and include the overnight move; the window Prev. 16:00–09:30 is the overnight gap from the preceding close to the announcement-day open; and the window 09:30–14:00 covers the announcement day from the open to the release. The three panels report all FOMC announcement days (All), the announcements before COVID (Pre-COVID, until 2020/03), and the announcements after it (Post-COVID, from 2020/04). The column SE reports the Newey–West standard errors. The column N reports the number of announcement days with available data in each window. The sample period is from 2016/09 to 2023/12.

All							
Window	$\Delta r_{t,t+h}$			$\Delta \log \text{VIX}_{t,t+h}$			N
	Mean	SE	SD	Mean	SE	SD	
Prev. 09:30–14:00	0.18	(0.12)	0.84	−0.71	(0.95)	6.53	59
Prev. 14:00–14:00	0.21**	(0.10)	0.67	−1.44	(0.88)	5.70	59
Prev. 16:00–09:30	0.19**	(0.07)	0.47	−0.85**	(0.39)	3.55	59
09:30–14:00	−0.03	(0.07)	0.54	−0.45	(0.54)	4.10	59
Pre-COVID							
Window	$\Delta r_{t,t+h}$			$\Delta \log \text{VIX}_{t,t+h}$			N
	Mean	SE	SD	Mean	SE	SD	
Prev. 09:30–14:00	0.09	(0.08)	0.48	−0.54	(1.14)	5.73	28
Prev. 14:00–14:00	0.08	(0.06)	0.36	−2.24*	(1.16)	4.69	28
Prev. 16:00–09:30	0.09***	(0.03)	0.22	−0.93**	(0.42)	2.61	28
09:30–14:00	−0.03	(0.07)	0.35	−1.14	(0.67)	4.07	28
Post-COVID							
Window	$\Delta r_{t,t+h}$			$\Delta \log \text{VIX}_{t,t+h}$			N
	Mean	SE	SD	Mean	SE	SD	
Prev. 09:30–14:00	0.25	(0.22)	1.08	−1.04	(1.56)	7.34	30
Prev. 14:00–14:00	0.37**	(0.17)	0.84	−1.11	(1.25)	6.17	30
Prev. 16:00–09:30	0.28**	(0.13)	0.62	−0.98	(0.70)	4.19	30
09:30–14:00	0.07	(0.07)	0.36	−0.33	(0.56)	3.02	30

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Table IV
Predicting Event Volatility on FOMC Announcement Days

This table reports the coefficients and R^2 of the predictive regression (33) of the realized event volatility on ATM implied volatility and spot volatility on the FOMC announcement days. The realized event volatility, $JV_{t,t+h}$, is measured by the difference between the realized volatility and the continuous volatility, as in equation (19), over the window $[t, t+h)$. The regressors are the ATM implied volatility, $\hat{I}_s(0, T)$, estimated from 0, 2, and 5DTE options, and the spot volatility, SV_s , described in Section 4. Both are measured at the information time s before 14:00 ET, the time of the announcement. The regression pools observations across days and intraday timestamps s , holding the $JV_{t,t+h}$ fixed at horizon $[t, t+h)$ within each day. Each variable is scaled to unit standard deviation and then expressed as a deviation from its timestamp- s mean, which is equivalent to absorbing one-way timestamp fixed effects, so the specification omits an intercept. The predictive regression is applied for two forecast windows: [14:00, 14:15) ET, the 15 minutes after the announcement, and [14:30, 15:00) ET, the half hour after the press conference starts at 14:30 ET. The parentheses in the table report standard errors clustered by day. The foot of each panel reports the adjusted R^2 (in %) from the bivariate regression with both factors and the number of minute-level observations N , and the $R^2_{\text{univ.}}$ (in %) from the univariate regression with each factor on the right side. The sample period is from 2016/09 to 2023/12.

DTE	Factor	JV(14:00–14:15)		JV(14:30–15:00)	
		Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
0	$\hat{I}_s(0, T)$	0.83*** (0.08)	78.12	0.84*** (0.09)	72.50
	SV_s	0.11* (0.06)	38.82	0.05 (0.07)	32.06
	R^2	78.40		72.12	
	N	13,440		13,440	
2	$\hat{I}_s(0, T)$	0.83*** (0.08)	67.96	0.89*** (0.11)	66.17
	SV_s	−0.01 (0.07)	37.05	−0.10 (0.08)	31.24
	R^2	67.38		66.00	
	N	13,440		13,440	
5	$\hat{I}_s(0, T)$	0.67*** (0.11)	58.33	0.72*** (0.15)	55.34
	SV_s	0.13 (0.09)	38.82	0.03 (0.10)	32.06
	R^2	58.26		54.56	
	N	13,440		13,440	

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Table V
Summary Statistics of Realized Volatility without Scheduled Events

This table reports summary statistics for ten-minute realized volatility and jump volatility, across three intraday windows, on non-event, FOMC Minutes release, and FOMC announcement days. Non-event days are the days excluding FOMC announcement and Minutes days and the COVID period (2020/03), with the intraday period from 10:00 ET to 15:40 ET. For each day type, three intraday windows are defined. The first runs from the open at 10:00 ET to 15:40 ET on non-event days, from 14:15 ET to 15:40 ET on Minutes days (the post-release window), and from 15:00 ET to 15:40 ET on FOMC announcement days (the post-press-conference window). The second is the pre-close window, [15:40, 15:50) ET, and the third is the closing window, [15:50, 16:00) ET. RV is realized volatility, computed as the square root of the realized variance from five-second ES futures log returns. JV is jump volatility, measured by the square root of the difference between the realized variance and the continuous variance, as in equation (19). Both quantities are annualized based on the stochastic clock (described in Section 3.1) and reported in percent. The columns report the number of observations (N), the mean value (Mean), the standard deviation (SD), and the 5th, 50th, and 95th percentiles (5%, 50%, and 95%) of the factor values within each window. The sample is from 2016/09 to 2023/12.

Period	Variables	Time	N	Mean	SD	5%	50%	95%
Non-Event	RV	[10:00, 15:40)	58,668	48.38***	25.61	23.48	40.49	100.21
		[15:40, 15:50)	1,691	53.68***	28.52	27.78	44.10	109.83
		[15:50, 16:00)	1,691	63.47***	33.48	33.48	52.50	126.33
	JV	[10:00, 15:40)	58,668	19.04***	16.34	0.00	16.38	48.29
		[15:40, 15:50)	1,691	24.27***	16.81	0.00	21.17	55.14
		[15:50, 16:00)	1,691	35.93***	23.06	12.27	30.78	74.40
Minutes	RV	[14:15, 15:40)	522	50.78***	24.59	25.48	42.75	101.31
		[15:40, 15:50)	58	54.91***	25.09	29.91	49.40	96.22
		[15:50, 16:00)	58	63.99***	39.25	34.28	54.37	120.10
	JV	[14:15, 15:40)	522	19.22***	15.46	0.00	17.22	46.95
		[15:40, 15:50)	58	25.92***	20.33	0.00	23.52	51.61
		[15:50, 16:00)	58	37.26***	33.63	16.74	29.54	66.97
FOMC	RV	[15:00, 15:40)	236	82.81***	54.16	33.41	63.79	207.88
		[15:40, 15:50)	59	75.23***	50.89	31.44	59.81	149.88
		[15:50, 16:00)	59	80.33***	50.09	35.89	62.68	158.16
	JV	[15:00, 15:40)	236	23.07***	23.69	0.00	20.46	66.56
		[15:40, 15:50)	59	28.17***	27.33	0.00	25.16	59.08
		[15:50, 16:00)	59	39.11***	28.76	0.00	34.96	113.78

Table VI
Summary Statistics of Expected Discontinuity Factors without Scheduled Events

This table reports summary statistics for the one-minute expected discontinuity factors, the magnitude, the net discontinuity, and the positive and negative discontinuities, estimated from 0–5DTE options on the days and intraday horizons without scheduled events. Non-event days are the days excluding FOMC announcement and Minutes days and the COVID period (2020/03), with the intraday period from 10:00 ET to 15:40 ET. The FOMC Minutes release and announcement days are the subsample of the scheduled-event days, but restricted to the intraday horizons after the events, from 14:15 ET to 15:40 ET for Minutes days and from 15:00 ET to 15:40 ET for FOMC announcement days. $\widehat{\mathcal{M}}^{N,J}$ is the implied discontinuity magnitude, $\widehat{\mathcal{J}}^N$ is the implied net discontinuity, and $\widehat{\mathcal{J}}^{N+}$ and $\widehat{\mathcal{J}}^{N-}$ are the implied upside and downside jumps, respectively. The estimation of the discontinuity factors is described in Section 4. The DTE column shows the days to expiration of the options used to estimate the factors. For each factor, the columns report the number of observations (N), the mean value (Mean), and the standard deviation (SD). The sample period is from 2016/09 to 2023/12.

Period	DTE	N	$\widehat{\mathcal{M}}^{N,J}$		$\widehat{\mathcal{J}}^N$		$\widehat{\mathcal{J}}^{N+}$		$\widehat{\mathcal{J}}^{N-}$	
			Mean	SD	Mean	SD	Mean	SD	Mean	SD
Non-Event	0	400,969	0.99***	1.19	−0.14***	0.37	0.42***	0.58	−0.57***	0.66
	1	407,763	0.38***	0.41	−0.01***	0.14	0.18***	0.21	−0.20***	0.23
	2	385,178	0.29***	0.30	0.01***	0.10	0.15***	0.16	−0.14***	0.16
	3	394,698	0.24***	0.25	0.02***	0.11	0.13***	0.13	−0.11***	0.13
	4	409,802	0.21***	0.20	0.03***	0.09	0.12***	0.11	−0.09***	0.11
	5	385,714	0.19***	0.19	0.04***	0.09	0.12***	0.11	−0.08***	0.10
Minutes	0	4,758	1.51***	1.53	−0.25***	0.56	0.63***	0.80	−0.88***	0.83
	2	4,590	0.25***	0.33	0.03***	0.08	0.14***	0.18	−0.11***	0.16
	5	4,760	0.17***	0.20	0.04***	0.09	0.10***	0.11	−0.06***	0.10
FOMC	0	2,240	1.83***	1.67	−0.51***	0.80	0.66***	0.90	−1.17***	0.95
	2	2,240	0.12***	0.20	−0.06***	0.16	0.03***	0.14	−0.09***	0.11
	5	2,240	0.08***	0.15	−0.02**	0.14	0.03***	0.11	−0.05***	0.08

Table VII
Predicting Closing Discontinuity Volatility on Non-Event Days

This table reports the coefficients and R^2 of the predictive regression (34) of the realized discontinuity volatility on expected discontinuity magnitude and spot volatility on non-event trading days. The sample is split into two groups: all non-event days and the normal days. Non-event days are the days excluding FOMC announcement and Minutes days and the COVID period (2020/03), with the intraday period from 10:00 ET to 15:40 ET. Normal days are the non-event days whose daily realized variance falls between the 25th and 75th percentiles. The realized jump volatility, $JV_{t,t+h}$, is measured by the difference between the realized volatility and the continuous volatility, as in equation (19), over the window $[t, t+h)$. The regressors are the expected discontinuity magnitude, $\widehat{\mathcal{M}}_{N,s}^{N,J}$, and the spot volatility, SV_s , measured at the information time s before 15:40 ET. The estimation of the factors is described in Section 4. The regression pools observations across days and intraday timestamps s , holding the $JV_{t,t+h}$ fixed at horizons [15:40, 15:50) ET and [15:50, 16:00) ET within each day. Each variable is scaled to unit standard deviation and then expressed as a deviation from its timestamp- s mean, which is equivalent to absorbing one-way timestamp fixed effects, so the specification omits an intercept. The table reports predictive results using factors estimated from 0-5DTE options, respectively. The parentheses in the table report standard errors clustered by day. The bottom of each panel reports the adjusted R^2 (in %) from the bivariate regression with both factors and the number of minute-level observations N , and the $R^2_{\text{univ.}}$ (in %) from the univariate regression with each factor on the right side. The sample period is from 2016/09 to 2023/12.

DTE	Factor	Non-Event				Normal			
		JV(15:40–15:50)		JV(15:50–16:00)		JV(15:40–15:50)		JV(15:50–16:00)	
		Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
0	$\widehat{\mathcal{M}}_s^{N,J}$	0.25*** (0.02)	8.48	0.29*** (0.03)	12.13	0.19*** (0.04)	1.54	0.35*** (0.03)	4.55
	SV_s	0.46*** (0.04)	24.19	0.57*** (0.05)	36.50	0.18*** (0.04)	1.95	0.39*** (0.03)	10.44
	R^2	28.72		42.88		4.23		18.49	
	N	399,267		399,267		196,643		196,643	
1	$\widehat{\mathcal{M}}_s^{N,J}$	0.16*** (0.02)	4.78	0.18*** (0.03)	6.44	0.08** (0.03)	0.54	0.22*** (0.05)	3.88
	SV_s	0.48*** (0.03)	25.11	0.56*** (0.06)	34.04	0.15*** (0.03)	2.11	0.38*** (0.03)	13.44
	R^2	27.40		37.15		2.58		17.77	
	N	406,812		406,812		211,078		211,078	
2	$\widehat{\mathcal{M}}_s^{N,J}$	0.12*** (0.03)	1.84	0.24*** (0.03)	6.65	0.03 (0.03)	0.02	0.20*** (0.05)	2.09
	SV_s	0.51*** (0.04)	25.93	0.60*** (0.05)	37.33	0.15*** (0.04)	2.15	0.39*** (0.03)	12.65
	R^2	27.13		42.59		2.09		16.13	
	N	384,305		384,305		192,663		192,663	
3	$\widehat{\mathcal{M}}_s^{N,J}$	0.14*** (0.03)	1.95	0.24*** (0.03)	5.51	0.04 (0.03)	0.03	0.25*** (0.04)	3.97
	SV_s	0.49*** (0.04)	23.92	0.61*** (0.05)	36.45	0.16*** (0.04)	2.40	0.40*** (0.03)	13.28
	R^2	25.70		41.71		2.37		19.07	
	N	392,489		392,489		199,483		199,483	
4	$\widehat{\mathcal{M}}_s^{N,J}$	0.11*** (0.02)	1.13	0.23*** (0.03)	5.13	0.03 (0.03)	0.04	0.24*** (0.04)	3.82
	SV_s	0.50*** (0.03)	24.96	0.62*** (0.05)	38.30	0.14*** (0.03)	1.91	0.39*** (0.03)	13.30
	R^2	26.01		43.35		1.86		18.48	
	N	408,516		408,516		209,713		209,713	
5	$\widehat{\mathcal{M}}_s^{N,J}$	0.11*** (0.03)	0.40	0.27*** (0.03)	4.13	0.02 (0.03)	0.00	0.23*** (0.04)	2.29
	SV_s	0.50*** (0.04)	24.03	0.63*** (0.06)	35.72	0.13*** (0.04)	1.62	0.39*** (0.04)	11.00
	R^2	25.15		42.39		1.48		15.70	
	N	382,941		382,941		189,253		189,253	

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Table VIII
Predicting Closing Discontinuity Volatility Using Uncertainty Increments on Non-Event Days

This table reports the coefficients and R^2 of the predictive regression (35) of the realized closing discontinuity volatility on the increment of the expected discontinuity magnitude, the level of the expected magnitude at a reference time, and spot volatility on days without scheduled events. The realized jump volatility, $JV_{t,t+h}$, is measured by the difference between the realized volatility and the continuous volatility, as in equation (19), over the closing window $[t, t+h) = [15:50, 16:00)$ ET. The regressors include the increment of the expected discontinuity magnitude, $\Delta\widehat{\mathcal{M}}_s^{N,J} = \widehat{\mathcal{M}}_s^{N,J} - \widehat{\mathcal{M}}_{\text{ref}}^{N,J}$, which is the change in the expected magnitude from the reference time to the sampling time s . The columns $\Delta(s_1)$, $\Delta(s_2)$, and $\Delta(s_3)$ correspond to the accumulation windows 10:00–12:00 ET, 12:00–14:00 ET, and 14:00–15:40 ET; each sets the reference time at the start of its window and the sampling time s at its end. The controls are the level of the expected discontinuity magnitude at the reference time, $\widehat{\mathcal{M}}_{\text{ref}}^{N,J}$, and the spot volatility, SV_s , described in Section 4. The factors are estimated from 0DTE, 1DTE, and 2DTE options, as the column groups indicate. Each column is a separate time-series regression across days, holding the forecast window fixed, with all variables scaled to unit standard deviation. The top panel uses the non-event days, the days excluding FOMC announcement and Minutes days and the COVID period (2020/03); the bottom panel uses the normal days, the non-event days whose daily realized variance falls between the 25th and 75th percentiles. The parentheses in the table report the Newey–West standard errors. The foot of each panel reports the R^2 (in %) and the number of days N . The sample period is from 2016/09 to 2023/12.

Non-Event									
Factor	0DTE			1DTE			2DTE		
	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$
α	0.50*** (0.08)	0.52*** (0.08)	0.55*** (0.07)	0.41*** (0.10)	0.42*** (0.10)	0.49*** (0.09)	0.24 (0.17)	0.24 (0.17)	0.36** (0.15)
$\Delta\widehat{\mathcal{M}}_s^{N,J}$	0.23*** (0.03)	0.27*** (0.03)	0.30*** (0.03)	0.14*** (0.03)	0.18*** (0.03)	0.22*** (0.03)	0.14*** (0.03)	0.26*** (0.06)	0.23*** (0.03)
$\widehat{\mathcal{M}}_{\text{ref}}^{N,J}$	0.15*** (0.03)	0.20*** (0.03)	0.14*** (0.03)	0.17** (0.08)	0.15*** (0.04)	0.15*** (0.03)	0.19** (0.08)	0.17*** (0.05)	0.21*** (0.05)
SV_s	0.54*** (0.05)	0.53*** (0.05)	0.58*** (0.05)	0.52*** (0.05)	0.56*** (0.06)	0.63*** (0.07)	0.59*** (0.08)	0.62*** (0.08)	0.65*** (0.09)
R^2	42.93	45.08	46.58	36.68	38.37	37.34	42.78	44.43	42.74
N	1,172	1,172	1,172	1,193	1,193	1,193	1,127	1,127	1,127
Normal									
Factor	0DTE			1DTE			2DTE		
	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$
α	1.08*** (0.14)	1.02*** (0.15)	1.04*** (0.12)	1.14*** (0.14)	1.00*** (0.14)	1.02*** (0.17)	1.19*** (0.14)	1.14*** (0.15)	1.23*** (0.16)
$\Delta\widehat{\mathcal{M}}_s^{N,J}$	0.22*** (0.05)	0.40*** (0.04)	0.38*** (0.04)	0.15*** (0.06)	0.24*** (0.03)	0.23*** (0.04)	0.18*** (0.05)	0.27*** (0.04)	0.27*** (0.05)
$\widehat{\mathcal{M}}_{\text{ref}}^{N,J}$	0.21*** (0.04)	0.24*** (0.05)	0.20*** (0.04)	0.13** (0.06)	0.14** (0.07)	0.16*** (0.06)	0.11* (0.06)	0.15** (0.06)	0.20*** (0.06)
SV_s	0.35*** (0.04)	0.36*** (0.05)	0.40*** (0.05)	0.37*** (0.05)	0.41*** (0.05)	0.50*** (0.07)	0.38*** (0.05)	0.41*** (0.06)	0.44*** (0.07)
R^2	18.88	21.03	26.08	17.05	18.54	21.26	15.92	16.41	17.04
N	577	577	577	619	619	619	565	565	565

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Table IX
Predicting Closing Discontinuity Volatility during Post-Event Horizons on Event Days

This table reports the coefficients and R^2 of the predictive regression (34) of the realized discontinuity volatility on expected discontinuity magnitude and spot volatility during the post-event horizons on FOMC Minutes release and announcement days. The post-event windows are from 14:15 ET to 15:40 ET for Minutes days and from 15:00 ET to 15:40 ET for FOMC announcement days. The realized jump volatility, $JV_{t,t+h}$, is measured by the difference between the realized volatility and the continuous volatility, as in equation (19), over the window $[t, t+h)$. The regressors are the expected discontinuity magnitude, $\widehat{\mathcal{M}}_{N,s}^{N,J}$, and the spot volatility, SV_s , measured at the information time s before 15:40 ET. The estimation of the factors is described in Section 4. The regression pools observations across days and intraday timestamps s , holding the $JV_{t,t+h}$ fixed at horizons [15:40, 15:50) ET and [15:50, 16:00) ET within each day. Each variable is scaled to unit standard deviation and then expressed as a deviation from its timestamp- s mean, which is equivalent to absorbing one-way timestamp fixed effects, so the specification omits an intercept. The table reports predictive results using factors estimated from 0, 2, and 5DTE options, respectively. The parentheses in the table report standard errors clustered by day. The bottom of each panel reports the adjusted R^2 (in %) from the bivariate regression with both factors and the number of minute-level observations N , and the $R^2_{\text{univ.}}$ (in %) from the univariate regression with each factor on the right side. The sample period is from 2016/09 to 2023/12.

DTE	Factor	Minutes				FOMC			
		JV(15:40–15:50)		JV(15:50–16:00)		JV(15:40–15:50)		JV(15:50–16:00)	
		Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
0	$\widehat{\mathcal{M}}_s^{N,J}$	0.32*** (0.11)	11.25	0.46** (0.20)	22.62	0.12 (0.09)	1.21	0.10 (0.09)	0.57
	SV_s	0.43*** (0.11)	20.41	0.62*** (0.18)	42.29	0.19 (0.23)	3.32	0.45** (0.21)	20.08
	R^2	28.48		60.35		2.92		19.48	
	N	4,814		4,814		2,296		2,296	
	$\widehat{\mathcal{M}}_s^{N,J}$	0.26 (0.26)	4.69	0.02 (0.12)	0.57	0.04 (0.07)	0.18	0.22* (0.12)	0.02
2	SV_s	0.33*** (0.12)	8.42	0.70*** (0.12)	48.23	0.21 (0.25)	3.82	0.54** (0.22)	20.31
	R^2	13.44		47.27		2.16		23.05	
	N	4,644		4,644		2,296		2,296	
	$\widehat{\mathcal{M}}_s^{N,J}$	0.34* (0.19)	5.49	0.35 (0.24)	4.19	−0.09 (0.10)	2.26	0.23* (0.12)	0.03
	SV_s	0.52*** (0.14)	20.44	0.72*** (0.26)	42.30	0.15 (0.25)	3.32	0.54** (0.23)	20.08
5	R^2	29.93		52.80		2.28		22.97	
	N	4,816		4,816		2,296		2,296	

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.

Appendix A. Derivations for Days with a Scheduled Event

Under the near-Gaussian shape of Assumption 1, the limits are taken as $T \downarrow 0$ and $\varepsilon \downarrow 0$, conditional on $\mathcal{F}_t$. Since $T_A \in (t, t+T)$ is deterministic and known at time t , the jump measure has the predictable fixed-time compensator stated in (2). There is no arrival-intensity term; conditional on the pre-event information, the only remaining uncertainty is Z_{T_A} . Conditional independence of Z_{T_A} and the diffusion paths, given $\mathcal{F}_t$, allows all event moments below to be evaluated under the time- t forecast law $\Pi_{A,t}$. For $n \geq 1$, let $c_{n,t}^A$ denote the n th conditional cumulant of Z_{T_A} under $\Pi_{A,t}$, finite at every order and bounded as in Assumption 1. Set $\sigma_{A,t}^2 = c_{2,t}^A$ and, for $n \geq 3$, define the standardized cumulants by $s_{n,t}^A := c_{n,t}^A / \sigma_{A,t}^n$.

Define the event log return using the martingale normalization

$$X_A := Z_{T_A} - \log \mathbb{E}_t^{\mathbb{Q}}[e^{Z_{T_A}}]. \quad (\text{A1})$$

Then $\mathbb{E}^{\mathbb{Q}}[e^{X_A} \mid \mathcal{F}_t] = 1$. The normalizing term is $\mathcal{F}_t$ -measurable, and hence

$$\text{Var}^{\mathbb{Q}}(X_A \mid \mathcal{F}_t) = c_{2,t}^A. \quad (\text{A2})$$

The background log-return component is

$$Y_T := \int_t^{t+T} \sigma_s dW_s - \frac{1}{2} \int_t^{t+T} \sigma_s^2 ds + \int_t^{t+T} \int x \tilde{\mu}(ds, dx) - \int_t^{t+T} \int (e^x - 1 - x) \nu_s(dx) ds. \quad (\text{A3})$$

Thus

$$R_{t,t+T} = X_A + Y_T. \quad (\text{A4})$$

By conditional independence, X_A and Y_T are independent conditional on $\mathcal{F}_t$. The variance process in (1) implies

$$\begin{aligned} \mathbb{E}^{\mathbb{Q}}[Y_T \mid \mathcal{F}_t] &= -\frac{1}{2} \sigma_t^2 T + o(T) \\ \text{Var}^{\mathbb{Q}}(Y_T \mid \mathcal{F}_t) &= \mathbb{E}^{\mathbb{Q}} \left[\int_t^{t+T} \sigma_s^2 ds \right] + \mathbb{E}^{\mathbb{Q}} \left[\int_t^{t+T} \int x^2 \nu_s(dx) ds \right] + o(T) \\ &= \sigma_t^2 T + \underbrace{T \int x^2 \nu_t(dx)}_{\sigma_{N,t}^2 T} + o(T). \end{aligned} \quad (\text{A5})$$

Consequently,

$$\text{Var}^{\mathbb{Q}}(R_{t,t+T} \mid \mathcal{F}_t) = c_{2,t}^A + \sigma_t^2 T + \sigma_{N,t}^2 T + o(T). \quad (\text{A6})$$

Let $c_{n,t}(T)$ denote the n th conditional cumulant of $R_{t,t+T}$. Equation (A6) gives

$$c_{2,t}(T) = c_{2,t}^A + \sigma_t^2 T + \sigma_{N,t}^2 T + o(T). \quad (\text{A7})$$

For $n \geq 3$, conditional independence and additivity of cumulants imply

$$c_{n,t}(T) = c_{n,t}^A + O(T), \quad n \geq 3, \quad (\text{A8})$$

because the non-announcement component is accumulated over an interval of length T .

Let $\Sigma_I(M, T) := I(m, T; \sigma_t) \sqrt{T}$ denote the total implied volatility as a function of the log-moneyness $M = \log(F_t/K) = \sigma_t \sqrt{T} m$. Applying Backus et al. (1997) to the $\mathcal{F}_t$ -conditional distribution of $R_{t,t+T}$ with cumulants $c_{n,t}(T)$ gives

$$\begin{aligned} \Sigma_I(M, T) \cong \sqrt{c_{2,t}(T)} & \left[1 - \frac{1}{3!} \frac{c_{3,t}(T)}{c_{2,t}(T)^{3/2}} d \right. \\ & \left. - \frac{1}{4!} \frac{c_{4,t}(T)}{c_{2,t}(T)^2} (1 - d^2) \right], \quad d = \frac{M + c_{2,t}(T)/2}{\sqrt{c_{2,t}(T)}}. \end{aligned} \quad (\text{A9})$$

Since

$$\begin{aligned} \frac{c_{3,t}(T)}{c_{2,t}(T)^{3/2}} &= \frac{c_{3,t}^A + O(T)}{\sigma_{A,t}^3 (1 + O(T))^{3/2}} = s_{3,t}^A + O(T), \\ \frac{c_{4,t}(T)}{c_{2,t}(T)^2} &= \frac{c_{4,t}^A + O(T)}{\sigma_{A,t}^4 (1 + O(T))^2} = s_{4,t}^A + O(T) = O(\varepsilon^2 \sigma_{A,t}^2) + O(T), \end{aligned} \quad (\text{A10})$$

equation (A9) can be rewritten as

$$\Sigma_I(M, T) \cong \sqrt{c_{2,t}(T)} \left(1 - \frac{1}{3!} s_{3,t}^A d \right) + O(\varepsilon^2 \sigma_{A,t}^2) + O(T). \quad (\text{A11})$$

At the money, $M = 0$ and $d = \frac{1}{2} \sqrt{c_{2,t}(T)}$, so equation (A11) is

$$I(0, T; \sigma_t) \sqrt{T} \cong \sqrt{c_{2,t}(T)} - \frac{1}{12} s_{3,t}^A \sigma_{A,t}^2 + O(\varepsilon^2 \sigma_{A,t}^2) + O(T). \quad (\text{A12})$$

Squaring both sides gives

$$\begin{aligned} I(0, T; \sigma_t)^2 T &\cong c_{2,t}(T) + O(\varepsilon \sigma_{A,t}^3, \varepsilon^2, \dots) + O(T) \\ &= \sigma_{A,t}^2 + \sigma_t^2 T + \sigma_{N,t}^2 T + O(\varepsilon \sigma_{A,t}^3, \varepsilon^2 \sigma_{A,t}^2, \dots) + O(T). \end{aligned} \quad (\text{A13})$$

The $O(T)$ term collects the contributions of the higher-order cumulants from the background components, which are of order $O(T)$ and can be ignored because the third and higher cumulants of the background component are small relative to its variance, which is itself retained explicitly. At the Gaussian benchmark, $\varepsilon = 0$, every shape term vanishes and the level reduces to

$$I(0, T; \sigma_t)^2 T = \sigma_{A,t}^2 + \sigma_t^2 T + \sigma_{N,t}^2 T + O(T). \quad (\text{A14})$$

Differentiating (A11) with respect to M , with $\partial d / \partial M = 1 / \sqrt{c_{2,t}(T)}$, and evaluating at $M = 0$ gives the ATM slope of the implied volatility smile,

$$\begin{aligned} \left. \frac{\partial \Sigma_I}{\partial M} \right|_{M=0} &\cong -\frac{1}{6} s_{3,t}^A + O(\varepsilon^2 \sigma_{A,t}^2) + O(T) \\ &= -\frac{1}{6} s_{3,t}^A + O(\varepsilon^2) + O(T). \end{aligned} \quad (\text{A15})$$

Since $\Sigma_I(M, T) = I(m, T; \sigma_t) \sqrt{T}$ with $M = \sigma_t \sqrt{T} m$, the chain rule gives

$$\left. \frac{\partial I}{\partial m} \right|_{m=0} = \sigma_t \left. \frac{\partial \Sigma_I}{\partial M} \right|_{M=0} \cong -\frac{\sigma_t}{6} s_{3,t}^A + O(\varepsilon^2) + O(T), \quad (\text{A16})$$

which yields the slope equation (5). □

Appendix B. Derivations for Days without a Scheduled Event

This appendix proves Proposition 2 under the absolutely continuous compensator (8). The argument extends the expansion method of Medvedev and Scaillet (2007), who derive the result for a compound-Poisson jump process, with constant intensity in their Proposition 4 and with intensity and jump-size distribution driven by the spot volatility in their Proposition 5, to an arbitrary predictable compensator $\nu_t(dx)$.

Under pure diffusion, $\nu_t = 0$, ATM implied volatility converges to the spot volatility

as $T \downarrow 0$. Following [Medvedev and Scaillet \(2007\)](#), assume that the pure-diffusion implied-volatility surface admits a local Taylor expansion around $(M, T) = (0, 0)$. Evaluating this expansion along $M = \sigma_t m \sqrt{T}$ gives

$$\begin{aligned} I(m, T; \sigma_t) &= \sigma_t + \left(\sigma_t m \frac{\partial I(0, 0; \sigma_t)}{\partial M} \right) \sqrt{T} + \left(\frac{\sigma_t^2 m^2}{2} \frac{\partial^2 I(0, 0; \sigma_t)}{\partial M^2} + \frac{\partial I(0, 0; \sigma_t)}{\partial T} \right) T + o(T) \\ &= \sigma_t + I_1(m) \sqrt{T} + I_2(m) T + o(T), \end{aligned} \tag{B1}$$

in which the coefficient I_1 is a linear and I_2 a quadratic polynomial in m . The second line is therefore an expansion, not an added restriction: given a smooth surface, the Taylor series delivers it and pins down the polynomial form of its coefficients.

With discontinuities, the surface is no longer well behaved at a fixed strike, and [Medvedev and Scaillet \(2007\)](#) show that the normalized moneyness m is the parameterization that replaces M . The reason is the behavior of options away from the money. The diffusion moves the price on the scale $\sigma_t \sqrt{T}$, so an option at fixed $M \neq 0$ pays off only if a discontinuity carries the price across the strike, an event of probability of order T under a finite arrival rate; matching a price of that size forces the implied volatility at fixed M to diverge as $T \downarrow 0$, so the Taylor series in the first line of (B1) no longer exists. Measured at fixed m , in units of the diffusive scale, the smile instead converges to the spot volatility, and the series in the second line of (B1) keeps its form, with coefficient functions that are no longer polynomials. That the series keeps this form, with twice continuously differentiable coefficients uniformly on compact sets of m , is the regularity the proof maintains. [Medvedev and Scaillet \(2007\)](#) verify it under compound-Poisson jumps by computing the short-maturity asymptotics of option prices directly. Here it is maintained for the general compensator, together with the finite-activity condition, $\nu_t(\mathbb{R}) < \infty$, and the exponential integrability, which make the total arrival rate and the functionals $\kappa_J(\nu_t)$ and $\kappa_{J\pm}(\nu_t)$ finite.

B.1 The pricing equation

Write the call price under the dynamics (1) as $c(S, V, T)$, a function of the underlying price, the spot variance, and the maturity. It satisfies the partial integro-differential equation

$$-c_T + \frac{1}{2} V S^2 c_{SS} + \kappa_v (\theta_v - V) c_V + \frac{1}{2} \sigma_v^2 V c_{VV} + \rho \sigma_v V S c_{SV} + \int_{\mathbb{R}} [c(S e^x) - c(S) - S(e^x - 1) c_S] \nu_t(dx) = 0, \quad (\text{B2})$$

where subscripts denote partial derivatives and the compensating term $S(e^x - 1) c_S$ reflects the martingale property of the underlying. Passing to the log moneyness through $C(M, V, T) := c(K e^M, V, T)$, the derivatives map into $c_S = C_M/S$, $c_{SS} = (C_{MM} - C_M)/S^2$, and $c_{SV} = C_{MV}/S$, and the pricing equation becomes

$$\underbrace{-C_T + \frac{1}{2} V (C_{MM} - C_M) + \kappa_v (\theta_v - V) C_V + \frac{1}{2} \sigma_v^2 V C_{VV} + \rho \sigma_v V C_{MV}}_{=: \mathcal{D}} + \mathcal{J} = 0, \quad (\text{B3})$$

with the jump operator

$$\mathcal{J} := \int_{\mathbb{R}} [C(M + x, V, T) - C(M, V, T) - (e^x - 1) C_M(M, V, T)] \nu_t(dx). \quad (\text{B4})$$

Substitute $C(M, V, T) = C^{\text{BS}}(M, I(m, T; \sigma_t), T)$, where $C^{\text{BS}}(M, I, T) = K(e^M N(d_1) - N(d_2))$ with

$$d_{1,2} = \frac{M}{I\sqrt{T}} \pm \frac{I\sqrt{T}}{2} = \frac{m\sigma_t}{I} \pm \frac{I\sqrt{T}}{2}, \quad (\text{B5})$$

and the BSM vega is $C_I^{\text{BS}} = K e^M \sqrt{T} n(d_1)$, with the identity $e^M n(d_1) = n(d_2)$. The chain

rule, with $\partial/\partial V = (2\sigma_t)^{-1} \partial/\partial\sigma_t$, gives

$$\begin{aligned}
C_T &= C_T^{\text{BS}} + C_I^{\text{BS}} I_T, \\
C_M &= C_M^{\text{BS}} + C_I^{\text{BS}} I_M, \\
C_{MM} &= C_{MM}^{\text{BS}} + 2C_{MI}^{\text{BS}} I_M + C_{II}^{\text{BS}} I_M^2 + C_I^{\text{BS}} I_{MM}, \\
C_V &= \frac{C_I^{\text{BS}} I_\sigma}{2\sigma_t}, \\
C_{VV} &= \frac{C_{II}^{\text{BS}} I_\sigma^2 + C_I^{\text{BS}} I_{\sigma\sigma}}{4\sigma_t^2} - \frac{C_I^{\text{BS}} I_\sigma}{4\sigma_t^3}, \\
C_{MV} &= \frac{C_{MI}^{\text{BS}} I_\sigma + C_{II}^{\text{BS}} I_M I_\sigma + C_I^{\text{BS}} I_{M\sigma}}{2\sigma_t},
\end{aligned} \tag{B6}$$

and the BSM derivatives reduce to multiples of the vega,

$$\begin{aligned}
C_T^{\text{BS}} &= \frac{I}{2T} C_I^{\text{BS}}, \\
C_{MM}^{\text{BS}} - C_M^{\text{BS}} &= \frac{C_I^{\text{BS}}}{IT}, \\
C_{MI}^{\text{BS}} &= -\frac{d_2}{I\sqrt{T}} C_I^{\text{BS}}, \\
C_{II}^{\text{BS}} &= \frac{d_1 d_2}{I} C_I^{\text{BS}}, \\
C_M^{\text{BS}} &= \frac{N(d_1)}{\sqrt{T} n(d_1)} C_I^{\text{BS}}.
\end{aligned} \tag{B7}$$

Dividing each term by the vega and substituting (B6) and (B7) gives

$$\begin{aligned}
\frac{C_T}{C_I^{\text{BS}}} &= \frac{I}{2T} + I_T, \\
\frac{C_{MM} - C_M}{C_I^{\text{BS}}} &= \frac{1}{IT} - \frac{2d_2}{I\sqrt{T}} I_M + \frac{d_1 d_2}{I} I_M^2 + I_{MM} - I_M, \\
\frac{C_V}{C_I^{\text{BS}}} &= \frac{I_\sigma}{2\sigma_t}, \\
\frac{C_{VV}}{C_I^{\text{BS}}} &= \left(\frac{d_1 d_2}{I} I_\sigma^2 + I_{\sigma\sigma} \right) \frac{1}{4\sigma_t^2} - \frac{I_\sigma}{4\sigma_t^3}, \\
\frac{C_{MV}}{C_I^{\text{BS}}} &= \left(-\frac{d_2}{I\sqrt{T}} I_\sigma + \frac{d_1 d_2}{I} I_M I_\sigma + I_{M\sigma} \right) \frac{1}{2\sigma_t}.
\end{aligned} \tag{B8}$$

Dividing both sides of (B3) by the BSM vega, C_I^{BS} , gives

$$\frac{\mathcal{D}}{C_I^{\text{BS}}} + \frac{\mathcal{J}}{C_I^{\text{BS}}} = 0. \quad (\text{B9})$$

B.2 Expansion of the diffusion part

Substituting the BSM derivatives (B8) into the diffusion part of (B9) yields

$$\begin{aligned} \frac{\mathcal{D}}{C_I^{\text{BS}}} = & -I_T + \frac{\sigma_t^2 - I^2}{2IT} + \frac{\sigma_t^2}{2} \left(-\frac{2d_2}{I\sqrt{T}} I_M + \frac{d_1 d_2}{I} I_M^2 + I_{MM} - I_M \right) + \left(\frac{\kappa_v(\theta_v - \sigma_t^2)}{2\sigma_t} - \frac{\sigma_v^2}{8\sigma_t} \right) I_\sigma \\ & + \frac{\sigma_v^2}{8} \left(\frac{d_1 d_2}{I} I_\sigma^2 + I_{\sigma\sigma} \right) + \frac{\rho\sigma_v\sigma_t}{2} \left(-\frac{d_2}{I\sqrt{T}} I_\sigma + \frac{d_1 d_2}{I} I_M I_\sigma + I_{M\sigma} \right). \end{aligned} \quad (\text{B10})$$

Substituting the representation (B1) and its powers,

$$\begin{aligned} I &= \sigma_t + I_1\sqrt{T} + I_2 T + o(T), \\ I^2 &= \sigma_t^2 + 2\sigma_t I_1\sqrt{T} + (I_1^2 + 2\sigma_t I_2) T + o(T), \\ I^3 &= \sigma_t^3 + 3\sigma_t^2 I_1\sqrt{T} + O(T), \end{aligned} \quad (\text{B11})$$

the derivatives of the implied volatility at fixed M become derivatives of the coefficient functions with respect to $m = M/(\sigma_t\sqrt{T})$, using $\partial m/\partial T = -m/(2T)$, $\partial m/\partial M = 1/(\sigma_t\sqrt{T})$, and $\partial m/\partial \sigma_t = -m/\sigma_t$:

$$\begin{aligned} I_T &= \left(\frac{1}{2}I_1 - \frac{1}{2}mI_{1m} \right) \frac{1}{\sqrt{T}} + I_2 - \frac{1}{2}mI_{2m} + O(\sqrt{T}), \\ I_M &= \frac{1}{\sigma_t}I_{1m} + \frac{1}{\sigma_t}I_{2m}\sqrt{T} + O(T), \\ I_{MM} &= \frac{1}{\sigma_t^2}I_{1mm} \frac{1}{\sqrt{T}} + \frac{1}{\sigma_t^2}I_{2mm} + O(\sqrt{T}), \\ I_\sigma &= 1 + \left(I_{1\sigma} - \frac{1}{\sigma_t}mI_{1m} \right) \sqrt{T} + O(T), \\ I_{\sigma\sigma} &= 0 + O(\sqrt{T}), \\ I_{M\sigma} &= \frac{1}{\sigma_t}I_{1m\sigma} - \frac{1}{\sigma_t^2}I_{1m} - \frac{1}{\sigma_t^2}mI_{1mm} + O(\sqrt{T}), \end{aligned} \quad (\text{B12})$$

where the subscripts m and σ denote differentiation of the coefficient functions with respect to m and the spot volatility. Substituting these replacements into (B10), with $d_1 = m\sigma_t/I + I\sqrt{T}/2$ and $d_2 = m\sigma_t/I - I\sqrt{T}/2$, collects the diffusion part into the form

$$\frac{\mathcal{D}}{C_I^{\text{BS}}} = A(m) \frac{1}{\sqrt{T}} + B(m) + O(\sqrt{T}), \quad (\text{B13})$$

where

$$\begin{aligned} A(m) &= \frac{1}{2}I_{1mm} - \frac{1}{2}mI_{1m} - \frac{3}{2}I_1 - \frac{\rho\sigma_v}{2}m, \\ B(m) &= -2I_2 - \frac{1}{2}mI_{2m} + \frac{1}{2}I_{2mm} + \frac{I_1^2}{2\sigma_t} + \frac{2m}{\sigma_t}I_1I_{1m} + \frac{m^2}{2\sigma_t}I_{1m}^2 \\ &\quad + \frac{\kappa_v(\theta_v - \sigma_t^2)}{2\sigma_t} - \frac{\sigma_v^2}{8\sigma_t} + \frac{\sigma_v^2 m^2}{8\sigma_t} + \frac{\rho\sigma_v\sigma_t}{4} \\ &\quad + \frac{m\rho\sigma_v}{\sigma_t}I_1 + \frac{\rho\sigma_v}{2} \left(\frac{2m^2}{\sigma_t}I_{1m} + I_{1m\sigma} - \frac{I_{1m}}{\sigma_t} - \frac{m}{\sigma_t}I_{1mm} - mI_{1\sigma} \right). \end{aligned} \quad (\text{B14})$$

Under pure diffusion the jump operator is absent and the pricing equation reduces to $A(m)/\sqrt{T} + B(m) + O(\sqrt{T}) = 0$, so both coefficients vanish. With $A(m) = 0$, the Taylor structure of (B1) restricts I_1 to the linear class, and the unique linear solution is

$$I_1(m) = -\frac{\rho\sigma_v}{4}m, \quad (\text{B15})$$

which is Proposition 1 of Medvedev and Scaillet (2007).⁶ At the money $I_1(0) = 0$: absent discontinuities, the ATM implied volatility equals the spot volatility up to $O(T)$, the benchmark against which the wedge in the level (10) is measured. Setting $B(m) = 0$ determines I_2 ; its solution is not needed at the stated $O(T)$ accuracy and is omitted.

⁶The variance dynamics in equation (1) imply $d\sigma_t = (\cdot)dt + (\sigma_v/2)dW_t^v$, so the volatility-of-volatility parameter b of Medvedev and Scaillet (2007) equals $\sigma_v/2$ here, and their pure-diffusion coefficient $I_1 = -\rho b\theta/2$ reads $-\rho\sigma_v m/4$.

B.3 Expansion of the jump operator

The jump operator normalized by the vega splits into two pieces,

$$\frac{\mathcal{J}}{C_I^{\text{BS}}} = \underbrace{\frac{1}{C_I^{\text{BS}}} \int_{\mathbb{R}} [C(M+x) - C(M)] \nu_t(dx)}_{=: \mathcal{R}_1} - \underbrace{\kappa_J(\nu_t) \frac{C_M}{C_I^{\text{BS}}}}_{=: \mathcal{R}_2}, \quad (\text{B16})$$

where $\kappa_J(\nu_t)$ collects the compensating terms. Since the vega carries the factor $\sqrt{T} n(d_1)$, both pieces call on the expansions of the BSM arguments at fixed m . Let $u = (I - \sigma_t)/\sigma_t$, so that the expansion of $1/I$ as $T \downarrow 0$ is the geometric series in u :

$$\begin{aligned} \frac{1}{I} &= \frac{1}{\sigma_t} (1 - u + u^2 - u^3 + O(u^4)) \\ &= \frac{1}{\sigma_t} - \frac{I_1}{\sigma_t^2} \sqrt{T} + \left(\frac{I_1^2}{\sigma_t^3} - \frac{I_2}{\sigma_t^2} \right) T + o(T). \end{aligned} \quad (\text{B17})$$

The arguments d_1 and d_2 then expand as

$$\begin{aligned} d_1 &= \frac{m\sigma_t}{I} + \frac{I\sqrt{T}}{2} = m + \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} + o(\sqrt{T}), \\ d_2 &= \frac{m\sigma_t}{I} - \frac{I\sqrt{T}}{2} = m + \left(-\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} + o(\sqrt{T}), \end{aligned} \quad (\text{B18})$$

so the normal distribution function evaluates to

$$\begin{aligned} N(d_1) &= N(m) + n(m) \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} + o(\sqrt{T}), \\ N(d_2) &= N(m) + n(m) \left(-\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} + o(\sqrt{T}), \end{aligned} \quad (\text{B19})$$

and the density, using $n'(m) = -m n(m)$, to

$$\begin{aligned} n(d_1) &= n(m) + n'(m) \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} + o(\sqrt{T}) \\ &= n(m) \left(1 - m \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} \right) + o(\sqrt{T}), \end{aligned} \quad (\text{B20})$$

whose reciprocal is

$$\begin{aligned}\frac{1}{n(d_1)} &= \frac{1}{n(m)} \frac{1}{1 - m\left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t}\right)\sqrt{T} + o(\sqrt{T})} \\ &= \frac{1}{n(m)} \left(1 + m\left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t}\right)\sqrt{T} + o(\sqrt{T})\right).\end{aligned}\tag{B21}$$

B.3.1 Expansion of $\mathcal{R}_1$

Write the shifted call price in $\mathcal{R}_1$ under the BSM formula as

$$\begin{aligned}C(M+x) &= K \left(e^{M+x} \underbrace{N\left(\frac{M+x}{I\sqrt{T}} + \frac{1}{2}I\sqrt{T}\right)}_{=: N_1(x)} - \underbrace{N\left(\frac{M+x}{I\sqrt{T}} - \frac{1}{2}I\sqrt{T}\right)}_{=: N_2(x)} \right) \\ &= K[e^{M+x}N_1(x) - N_2(x)],\end{aligned}\tag{B22}$$

so that

$$\begin{aligned}\mathcal{R}_1 &= \frac{1}{C_I^{\text{BS}}} \int_{\mathbb{R}} [C(M+x) - C(M)] \nu_t(dx) \\ &= \frac{1}{\sqrt{T}n(d_1)} \left(\int_{\mathbb{R}} [e^x N_1(x) - N(d_1)] \nu_t(dx) - e^{-M} \int_{\mathbb{R}} [N_2(x) - N(d_2)] \nu_t(dx) \right).\end{aligned}\tag{B23}$$

At fixed m , the moneyness enters the arguments of N_1 and N_2 only through $M/(I\sqrt{T}) = m\sigma_t/I$, which is finite. For any jump $x \neq 0$, the term $\frac{x}{I\sqrt{T}} \rightarrow \pm\infty$ as $T \downarrow 0$. Hence, the CDF terms converge to 0 or 1:

$$\lim_{T \downarrow 0} N_1(x) = \mathbf{1}_{\{x>0\}}, \quad \lim_{T \downarrow 0} N_2(x) = \mathbf{1}_{\{x>0\}}.\tag{B24}$$

The bounds $e^x N_1(x) \leq e^x$ and $N_2(x) \leq 1$ justify moving these limits inside the integrals of equation (B23) by dominated convergence. The compound-Poisson structure of [Medvedev and Scaillet \(2007\)](#) plays no role here: the limits are properties of the BSM functionals alone, and the interchange needs only the finite arrival rate $\nu_t(\mathbb{R}) < \infty$ and the exponential integrability. Making the substitution, I separate positive from negative jumps and write

$\nu_{\pm} := \nu_t(\mathbb{R}_{\pm})$ for the arrival rates. On $\mathbb{R}_+$, where $N_1(x) \rightarrow 1$ and $N_2(x) \rightarrow 1$,

$$\begin{aligned}
\mathcal{R}_1^+ &= \frac{1}{\sqrt{T} n(d_1)} \left(\int_{\mathbb{R}_+} (e^x - 1) \nu_t(dx) + \int_{\mathbb{R}_+} [1 - N(d_1)] \nu_t(dx) - e^{-M} \int_{\mathbb{R}_+} [1 - N(d_2)] \nu_t(dx) \right) \\
&= \frac{1}{\sqrt{T} n(d_1)} \left(\kappa_{J+}(\nu_t) + [1 - N(d_1)] \nu_+ - e^{-M} [1 - N(d_2)] \nu_+ \right) \\
&= \frac{1}{\sqrt{T} n(d_1)} \left(\kappa_{J+}(\nu_t) + \left(1 - N(m) - n(m) \left(\frac{\sigma_t}{2} - \frac{m I_1}{\sigma_t} \right) \sqrt{T} \right) \nu_+ \right) \\
&\quad - \frac{1}{\sqrt{T} n(d_1)} e^{-M} \left(1 - N(m) + n(m) \left(\frac{\sigma_t}{2} + \frac{m I_1}{\sigma_t} \right) \sqrt{T} \right) \nu_+,
\end{aligned} \tag{B25}$$

by the expansions (B19). Substituting $e^{-M} = 1 - \sigma_t m \sqrt{T} + O(T)$ gives

$$\mathcal{R}_1^+ = \frac{\kappa_{J+}(\nu_t)}{n(d_1)} \frac{1}{\sqrt{T}} + \frac{\sigma_t}{n(d_1)} [m - n(m) - m N(m)] \nu_+ + O(\sqrt{T}). \tag{B26}$$

On $\mathbb{R}_-$, where $N_1(x) \rightarrow 0$ and $N_2(x) \rightarrow 0$, the same substitutions give

$$\begin{aligned}
\mathcal{R}_1^- &= \frac{1}{\sqrt{T} n(d_1)} \left(- \int_{\mathbb{R}_-} N(d_1) \nu_t(dx) + e^{-M} \int_{\mathbb{R}_-} N(d_2) \nu_t(dx) \right) \\
&= \frac{1}{\sqrt{T} n(d_1)} \left(- N(d_1) \nu_- + e^{-M} N(d_2) \nu_- \right) \\
&= - \frac{\sigma_t}{n(d_1)} [n(m) + m N(m)] \nu_- + O(\sqrt{T}).
\end{aligned} \tag{B27}$$

Summing the two sides,

$$\begin{aligned}
\mathcal{R}_1 &= \mathcal{R}_1^+ + \mathcal{R}_1^- \\
&= \frac{\kappa_{J+}(\nu_t)}{n(d_1)} \frac{1}{\sqrt{T}} + \frac{\sigma_t}{n(d_1)} [m - n(m) - m N(m)] \nu_+ - \frac{\sigma_t}{n(d_1)} [n(m) + m N(m)] \nu_- + O(\sqrt{T}) \\
&= \frac{\kappa_{J+}(\nu_t)}{n(d_1)} \frac{1}{\sqrt{T}} + \frac{\sigma_t}{n(d_1)} \left[m \nu_+ - (n(m) + m N(m)) (\nu_+ + \nu_-) \right] + O(\sqrt{T}),
\end{aligned} \tag{B28}$$

and substituting the expansion of $1/n(d_1)$ in (B21),

$$\begin{aligned}
\mathcal{R}_1 &= \frac{\kappa_{J+}(\nu_t)}{n(m)} \frac{1}{\sqrt{T}} \left(1 + m \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} \right) + \frac{\sigma_t}{n(m)} \left[m \nu_+ - (n(m) + mN(m))(\nu_+ + \nu_-) \right] + O(\sqrt{T}) \\
&= \underbrace{\frac{\kappa_{J+}(\nu_t)}{n(m)} \frac{1}{\sqrt{T}}}_{=: \hat{A}_1(m)} + \underbrace{\frac{\kappa_{J+}(\nu_t)}{n(m)} m \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) + \frac{\sigma_t}{n(m)} \left[m \nu_+ - (n(m) + mN(m))(\nu_+ + \nu_-) \right]}_{=: \hat{B}_1(m)} + O(\sqrt{T}).
\end{aligned} \tag{B29}$$

B.3.2 Expansion of $\mathcal{R}_2$

From the chain rule (B6) and the BSM formula,

$$\frac{C_M}{C_I^{\text{BS}}} = \frac{C_M^{\text{BS}}}{C_I^{\text{BS}}} + I_M = \frac{Ke^M N(d_1)}{Ke^M \sqrt{T} n(d_1)} + I_M = \frac{N(d_1)}{\sqrt{T} n(d_1)} + I_M. \tag{B30}$$

Substituting the expansions (B12), (B19), and (B21), and $I_M = I_{1m}/\sigma_t + O(\sqrt{T})$ in the last step,

$$\begin{aligned}
\mathcal{R}_2 &= \kappa_J(\nu_t) \left(\frac{N(d_1)}{\sqrt{T} n(d_1)} + I_M \right) \\
&= \frac{\kappa_J(\nu_t)}{n(m) \sqrt{T}} \left(N(m) + n(m) \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} \right) \left(1 + m \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) \sqrt{T} \right) + \kappa_J(\nu_t) I_M + O(\sqrt{T}) \\
&= \frac{\kappa_J(\nu_t)}{\sqrt{T}} \frac{N(m)}{n(m)} + \frac{\kappa_J(\nu_t)}{n(m)} \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) (n(m) + mN(m)) + \kappa_J(\nu_t) I_M + O(\sqrt{T}) \\
&= \underbrace{\kappa_J(\nu_t) \frac{N(m)}{n(m)} \frac{1}{\sqrt{T}}}_{=: \hat{A}_2(m)} + \underbrace{\frac{\kappa_J(\nu_t)}{n(m)} \left(\frac{\sigma_t}{2} - \frac{mI_1}{\sigma_t} \right) (n(m) + mN(m)) + \frac{\kappa_J(\nu_t)}{\sigma_t} I_{1m}}_{=: \hat{B}_2(m)} + O(\sqrt{T}).
\end{aligned} \tag{B31}$$

Finally, combining (B13), (B14), (B29), and (B31), the pricing equation (B3) normalized by the vega reads

$$0 = \frac{\mathcal{D}}{C_I^{\text{BS}}} + \mathcal{R}_1 - \mathcal{R}_2 = \left(A(m) + \hat{A}_1(m) - \hat{A}_2(m) \right) \frac{1}{\sqrt{T}} + \left(B(m) + \hat{B}_1(m) - \hat{B}_2(m) \right) + O(\sqrt{T}). \tag{B32}$$

The coefficient at order $T^{-1/2}$ depends on ν_t only through the two scalars $\kappa_J(\nu_t)$ and $\kappa_{J+}(\nu_t)$;

the arrival rates $\nu_{\pm}$, and any finer feature of the compensator, enter only the order-one coefficient, which governs I_2 .

B.4 The solution of the ODE for I_1

Setting the coefficient of $T^{-1/2}$ in (B32) to zero, $A(m) + \hat{A}_1(m) - \hat{A}_2(m) = 0$, gives the ordinary differential equation

$$\begin{aligned} 0 &= \frac{1}{2}I_{1mm} - \frac{1}{2}mI_{1m} - \frac{3}{2}I_1 - \frac{\rho\sigma_v}{2}m + \frac{\kappa_{J+}(\nu_t)}{n(m)} - \kappa_J(\nu_t)\frac{N(m)}{n(m)} \\ \Rightarrow \frac{1}{2}I_{1mm} - \frac{1}{2}mI_{1m} - \frac{3}{2}I_1 &= \frac{\rho\sigma_v}{2}m + \kappa_J(\nu_t)\frac{N(m)}{n(m)} - \frac{\kappa_{J+}(\nu_t)}{n(m)}. \end{aligned} \quad (\text{B33})$$

The solution assembles from three particular solutions, one per source term. For the leverage term, the pure-diffusion solution (B15) applies: $u(m) = -\frac{\rho\sigma_v}{4}m$ satisfies

$$\frac{1}{2}u_{mm} - \frac{1}{2}mu_m - \frac{3}{2}u = \frac{\rho\sigma_v}{2}m. \quad (\text{B34})$$

For the second term, let $u(m) = -\kappa_J(\nu_t)N(m)/n(m)$; using $n'(m) = -m n(m)$, its derivatives are

$$\begin{aligned} u_m &= -\kappa_J(\nu_t) \left(\frac{n(m)}{n(m)} - \frac{N(m)n'(m)}{n(m)^2} \right) = -\kappa_J(\nu_t) \left(1 + m \frac{N(m)}{n(m)} \right), \\ u_{mm} &= -\kappa_J(\nu_t) \left(m + \frac{N(m)}{n(m)} + m^2 \frac{N(m)}{n(m)} \right), \end{aligned} \quad (\text{B35})$$

so that

$$\frac{1}{2}u_{mm} - \frac{1}{2}mu_m - \frac{3}{2}u = \kappa_J(\nu_t) \frac{N(m)}{n(m)}. \quad (\text{B36})$$

For the third term, let $u(m) = \kappa_{J+}(\nu_t)/n(m)$, with derivatives

$$\begin{aligned} u_m &= m \frac{\kappa_{J+}(\nu_t)}{n(m)}, \\ u_{mm} &= \frac{\kappa_{J+}(\nu_t)}{n(m)} + m^2 \frac{\kappa_{J+}(\nu_t)}{n(m)}, \end{aligned} \quad (\text{B37})$$

so that

$$\frac{1}{2}u_{mm} - \frac{1}{2}mu_m - \frac{3}{2}u = -\frac{\kappa_{J+}(\nu_t)}{n(m)}. \quad (\text{B38})$$

Summing the three particular solutions shows that

$$I_1(m) = -\frac{\rho\sigma_v}{4}m - \kappa_J(\nu_t)\frac{N(m)}{n(m)} + \kappa_{J+}(\nu_t)\frac{1}{n(m)} \quad (\text{B39})$$

solves (B33), which is the $\sqrt{T}$ coefficient in equation (9).

At the money, $N(0)/n(0) = \sqrt{\pi/2}$, $1/n(0) = \sqrt{2\pi}$, and $\kappa_J = \kappa_{J+} + \kappa_{J-}$, so

$$I_1(0) = \sqrt{2\pi}\kappa_{J+}(\nu_t) - \sqrt{\frac{\pi}{2}}\kappa_J(\nu_t) = \sqrt{\frac{\pi}{2}}(\kappa_{J+}(\nu_t) - \kappa_{J-}(\nu_t)), \quad (\text{B40})$$

which is the level equation (10). For the slope, the uniformity of the representation (B1) on compact sets of m allows term-by-term differentiation, $\partial I/\partial m = I_{1m}(m)\sqrt{T} + O(T)$, and differentiating (B39), using the derivatives in (B35) and (B37), gives

$$\begin{aligned} I_{1m}(m) &= -\frac{\rho\sigma_v}{4} - \kappa_J(\nu_t) \left(1 + m \frac{N(m)}{n(m)}\right) + \kappa_{J+}(\nu_t) \frac{m}{n(m)}, \\ I_{1m}(0) &= -\frac{\rho\sigma_v}{4} - \kappa_J(\nu_t), \end{aligned} \quad (\text{B41})$$

which is the slope equation (11). □

Remark. For the compound-Poisson specification of Medvedev and Scaillet (2007), $\nu_t(dx) = \lambda F(dx)$ with relative jump size $\Delta J := e^x - 1$, the functionals reduce to $\kappa_J(\nu_t) = \lambda \mathbb{E}[\Delta J]$ and $\kappa_{J+}(\nu_t) = \lambda \mathbb{E}[\Delta J^+]$, and equation (B39) coincides with their Proposition 4. Besides, their Proposition 5, which lets the intensity and the jump-size distribution vary with the spot volatility, leaves I_1 unchanged in the same way. The generalization to the compensator (8) therefore alters the expansion only in interpretation: κ_J and κ_{J+} are conditional, time-varying functions of the current compensator rather than constants. The derivation shows why. The jump operator enters the $T^{-1/2}$ balance only through the compensator frozen at its time- t value, and any state variable that drives ν_t moves the implied volatility only through the $\sqrt{T}$ coefficient, so the associated generator terms are $O(1)$ after vega normalization and cannot disturb the balance that determines I_1 . The dynamics of the compensator, whether driven by time, by volatility, or by additional state variables, surface first in the $O(T)$ term I_2 . The identification of the expected jump magnitude in Section 2.2 is therefore robust to the specification of the arrival process and of the jump-size distribution.

Appendix C. The Leverage Contribution to the Smile Slope

This appendix estimates the leverage term $\rho\sigma_v/4$ in the expansion (11) with the option-based estimator of Chong and Todorov (2024). Under the dynamics of equation (1), the log spot variance satisfies $d \log \sigma_t^2 = (\sigma_v/\sigma_t) dW_t^v + (\cdots) dt$, so its diffusive covariation with the return identifies the leverage parameter,

$$\frac{d[S, \log \sigma^2]_t}{dt} = \rho\sigma_v. \quad (\text{C1})$$

I apply Chong and Todorov (2024)'s single-maturity estimator to the one-minute series of SV. Let x_{t_i} be the log ES futures price and $\hat{\sigma}_{t_i}^2 = \text{SV}_{t_i}^2$ the option-implied spot variance, and set the log-variance $\hat{V}_{t_i} = \log \hat{\sigma}_{t_i}^2$, with increments $\Delta_i x = x_{t_i} - x_{t_{i-1}}$ and $\Delta_i \hat{V} = \hat{V}_{t_i} - \hat{V}_{t_{i-1}}$. Applying their estimator over the whole sample gives

$$\widehat{\rho\sigma_v} = \frac{1}{N \Delta_n} \sum_i \tau_{v_n}(\Delta_i x) \tau_{v_n}(\Delta_i \hat{V}), \quad \tau_v(z) = z \mathbf{1}\{|z| \leq v\}, \quad (\text{C2})$$

where $\Delta_n = 1/\Omega$ is the one-minute step on the variance clock of Section 3.1 and N is the number of increments in the 10:00–15:40 window. Following Chong and Todorov (2024), the same truncation applies to both increments, with the threshold set from a four-day bipower variation of $\Delta_i \hat{V}$,

$$v_n = 3 \left(\frac{\pi}{8} \frac{1}{(k_n - 1)\Delta_n} \sum_{l=0}^3 \sum_{i=2}^{k_n} |\Delta_i \hat{V}_{t-l/252}| |\Delta_{i-1} \hat{V}_{t-l/252}| \right)^{0.49}, \quad (\text{C3})$$

with k_n the number of within-day increments. The option observation error requires no bias correction here, unlike for volatility of volatility, because it is orthogonal to the return and so introduces no bias.

The pooled estimate is $\widehat{\rho\sigma_v} = -0.61$, with a t -statistic of -10.6 . Since ρ and σ_v are structural and slowly varying, I use the single-maturity form of the estimator to obtain a constant estimate over the whole sample. The two-maturity version of Chong and Todorov (2024) further reduces the option-error bias, but the processed SV series is a single maturity, and their Monte Carlo shows the single-maturity estimator is only slightly less precise.

Figure A.I. Intraday Patterns of Realized Volatility with Alternative Truncation Parameters on FOMC Announcement Days

This figure plots the intraday patterns of one-minute realized volatility measures from 10:00 to 16:00 ET, averaged across FOMC announcement days, under alternative truncation parameters for the continuous volatility. The columns show the realized volatility, RV, the continuous volatility, CV, and the jump volatility, JV, estimated from five-second ES futures returns as described in Section 4.1. The truncation threshold of equation (18) multiplies the local volatility level, estimated over the preceding window of w minutes, by the constant τ . The top row varies the window, $w = 5, 10$, and 30 minutes, holding $\tau = 2.5$; the bottom row varies the multiple, $\tau = 2.0, 2.5$, and 3.0, holding $w = 30$ minutes, so the baseline choice of the paper, $\tau = 2.5$ and $w = 30$ minutes, appears in both rows. The RV panels do not depend on the truncation and are identical across the two rows, providing the common benchmark that CV and JV decompose. All variables are annualized based on the stochastic clock (described in Section 3.1) and expressed in percent. The sample period is from 2016/09 to 2023/12.

A16




Figure A.II. Intraday Patterns of SV Using Different Estimation Adjustments

This figure plots the intraday patterns of the spot volatility, SV, under different estimation adjustments. The estimate on a given minute can be invalid because there are not enough option quotes available or the estimator does not return a finite value. The no-adjustment series uses only the shortest available maturity, leaving the minute missing when its estimate is invalid. The fallback series starts from the shortest available maturity and, when the estimate is invalid, moves to the next longer maturity until a valid estimate is obtained. The two-maturity series averages the spot variances of the two shortest maturities with valid estimates, moving to the next longer maturity in the same way when an estimate is invalid, and takes the square root of the average. The left and right columns plot the sample across all trading days and across FOMC announcement days, respectively, from 10:00 to 15:40 ET. The vertical dashed lines in the FOMC panels mark 14:00 and 14:30 ET, the scheduled release and press conference times. The top row plots the average SV at each minute, and the bottom row plots the number of observations with a valid estimate at each minute. The sample period is from 2016/09 to 2023/12.




Figure A.III. Pre-FOMC Announcement Drift and VIX Change across One-Day Windows

The figure plots the cumulative returns of the ES futures and the cumulative change of the VIX over the FOMC announcement day. The top row shows the average cumulative log return of the ES futures, accumulated from 09:30 ET on the FOMC announcement day. The bottom row shows the average cumulative log change of the 30-day VIX from the same starting point, where the VIX is computed from SPX options at the one-minute frequency following the standard Cboe methodology, with maturities measured by the stochastic maturity of Section 3.1. The columns split the announcements into two samples: all announcements and announcements after COVID (from 2020/04). Shaded areas are pointwise 95% confidence bands across announcements, and the vertical dashed line marks the 14:00 ET announcement time. The sample period is from 2016/09 to 2023/12.




Figure A.IV. Pre-FOMC Announcement Drift and VIX Change under Calendar Maturity across Two-Day Windows

The figure plots the cumulative returns of the ES futures and the cumulative change of the VIX over the two-day window around FOMC announcements. The top row shows the average cumulative log return of the ES futures, accumulated from 09:30 ET of the trading day preceding the announcement. The bottom row shows the average cumulative log change of the 30-day VIX from the same starting point, where the VIX is computed from SPX options at the one-minute frequency following the standard Cboe methodology with calendar maturities. The columns split the announcements into two samples: all announcements and announcements after COVID (from 2020/04). Shaded areas are pointwise 95% confidence bands across announcements, and the vertical dashed line marks the 14:00 ET announcement time. The sample period is from 2016/09 to 2023/12.




Figure A.V. Pre-FOMC Announcement Drift and VIX Change under Calendar Maturity across One-Day Windows

The figure plots the cumulative returns of the ES futures and the cumulative change of the VIX over the FOMC announcement day. The top row shows the average cumulative log return of the ES futures, accumulated from 09:30 ET on the FOMC announcement day. The bottom row shows the average cumulative log change of the 30-day VIX from the same starting point, where the VIX is computed from SPX options at the one-minute frequency following the standard Cboe methodology with calendar maturities. The columns split the announcements into two samples: all announcements and announcements after COVID (from 2020/04). Shaded areas are pointwise 95% confidence bands across announcements, and the vertical dashed line marks the 14:00 ET announcement time. The sample period is from 2016/09 to 2023/12.




Figure A.VI. Intraday Patterns of ATM Implied Volatility under Different Maturity Measurements

The figure plots the intraday patterns of ATM implied volatility under three measures of maturity. The stochastic maturity, introduced in Section 3.1, counts the remaining time on a variance clock that weights each minute by its share of variance accumulation. The calendar maturity uses calendar time to expiration, counting 24 hours per day, and the business maturity uses the remaining trading hours, counting 6.5 hours per day from 9:30 to 16:00 ET. The three maturities share the same option prices, so the differences come only from the measure of maturity. The top row plots FOMC announcement days from 10:00 to 15:00 ET for 0DTE, 2DTE, and 5DTE options, with the vertical dashed lines marking 14:00 and 14:30 ET, the scheduled release and press conference times, respectively. The bottom two rows plot non-FOMC days, the days without FOMC announcements, from 10:00 to 15:40 ET for 0DTE to 5DTE options. All values are expressed in percent. The sample period is from 2016/09 to 2023/12.

A21




Table A.I
The Share of Maturities of Options Used in the Spot Volatility Estimation

This table reports the distribution of the option maturities used in the spot volatility estimation, described in Section 4.2, under two adjustments to handle invalid estimates. The estimate on a given minute can be invalid because there are not enough option quotes available or the estimator does not return a finite value. Each row in the table shows the percentage of the one-minute observations for which the estimation uses the given maturities, and the foot of each panel reports the number of one-minute observations. Panel A reports the fallback method, which starts from the shortest available maturity and, when the estimate is invalid, moves to the next longer maturity until a valid estimate is obtained. Panel B reports the two-maturity method, which averages the spot variances of the two shortest maturities with valid estimates, moving to the next longer maturity in the same way when an estimate is invalid. The sample period is from 2016/09 to 2023/12.

Panel A: Fallback (one contract)			
Maturity used		All Days	FOMC Days
0DTE		69.1	93.1
1DTE		29.1	5.4
2DTE		1.7	1.5
3DTE		0.2	0.1
≥ 4 DTE		0.0	0.0
Observations		627,930	20,118
Panel B: Two-maturity (two contracts)			
Short	Long	All Days	FOMC Days
0DTE	1DTE	38.6	26.8
0DTE	2DTE	29.9	62.9
0DTE	3DTE	0.5	3.3
1DTE	2DTE	14.0	3.7
1DTE	3DTE	14.7	1.7
2DTE	3DTE	0.6	1.5
2DTE	4DTE	1.0	0.0
Other		0.6	0.1
Observations		627,892	20,118

Table A.II
Summary Statistics of the Pre-Announcement Drift and VIX Change under Calendar Maturity

This table reports the mean and standard deviation of the pre-announcement cumulative log returns of the ES futures and the cumulative change of the VIX on FOMC announcement days. The $\Delta r_{t,t+h}$ is the cumulative log return of the ES futures from the start of the window t to its end time $t+h$. The $\Delta \log \text{VIX}_{t,t+h}$ is the cumulative log change of the 30-day VIX over the same window, in percent, where the VIX is computed from SPX options following the standard Cboe methodology, with calendar maturities. The windows Prev. 09:30–14:00 and Prev. 14:00–14:00 start at 09:30 ET and 14:00 ET of the trading day preceding the announcement, end at the 14:00 ET release, and include the overnight move; the window Prev. 16:00–09:30 is the overnight gap from the preceding close to the announcement-day open; and the window 09:30–14:00 covers the announcement day from the open to the release. The three panels report all FOMC announcement days (All), the announcements before COVID (Pre-COVID, until 2020/03), and the announcements after it (Post-COVID, from 2020/04). The column SE reports the Newey–West standard errors. The column N reports the number of announcement days with available data in each window. The sample period is from 2016/09 to 2023/12.

All							
Window	$\Delta r_{t,t+h}$			$\Delta \log \text{VIX}_{t,t+h}$			N
	Mean	SE	SD	Mean	SE	SD	
Prev. 09:30–14:00	0.18	(0.12)	0.84	−2.03**	(0.94)	6.46	59
Prev. 14:00–14:00	0.21**	(0.10)	0.67	−2.14**	(0.88)	5.68	59
Prev. 16:00–09:30	0.19**	(0.07)	0.47	−0.67*	(0.39)	3.54	59
09:30–14:00	−0.03	(0.07)	0.54	−1.03*	(0.54)	4.11	59
Pre-COVID							
Window	$\Delta r_{t,t+h}$			$\Delta \log \text{VIX}_{t,t+h}$			N
	Mean	SE	SD	Mean	SE	SD	
Prev. 09:30–14:00	0.09	(0.08)	0.48	−1.87	(1.14)	5.71	28
Prev. 14:00–14:00	0.08	(0.06)	0.36	−2.95**	(1.16)	4.67	28
Prev. 16:00–09:30	0.09***	(0.03)	0.22	−0.76*	(0.41)	2.60	28
09:30–14:00	−0.03	(0.07)	0.35	−1.72**	(0.68)	4.07	28
Post-COVID							
Window	$\Delta r_{t,t+h}$			$\Delta \log \text{VIX}_{t,t+h}$			N
	Mean	SE	SD	Mean	SE	SD	
Prev. 09:30–14:00	0.25	(0.22)	1.08	−2.37	(1.51)	7.20	30
Prev. 14:00–14:00	0.37**	(0.17)	0.84	−1.83	(1.23)	6.13	30
Prev. 16:00–09:30	0.28**	(0.13)	0.62	−0.80	(0.70)	4.18	30
09:30–14:00	0.07	(0.07)	0.36	−0.92	(0.55)	3.03	30

* denotes statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Table A.III
Predicting Closing Discontinuity Volatility Using Uncertainty Increments from Longer-Maturity Options on Non-Event Days

This table reports the coefficients and R^2 of the predictive regression (35) of the realized closing discontinuity volatility on the increment of the expected discontinuity magnitude, the level of the expected magnitude at a reference time, and spot volatility on days without scheduled events. The realized jump volatility, $JV_{t,t+h}$, is measured by the difference between the realized volatility and the continuous volatility, as in equation (19), over the closing window $[t, t+h) = [15:50, 16:00)$ ET. The regressors include the increment of the expected discontinuity magnitude, $\Delta\widehat{\mathcal{M}}_s^{N,J} = \widehat{\mathcal{M}}_s^{N,J} - \widehat{\mathcal{M}}_{\text{ref}}^{N,J}$, which is the change in the expected magnitude from the reference time to the sampling time s . The columns $\Delta(s_1)$, $\Delta(s_2)$, and $\Delta(s_3)$ correspond to the accumulation windows 10:00–12:00 ET, 12:00–14:00 ET, and 14:00–15:40 ET; each sets the reference time at the start of its window and the sampling time s at its end. The controls are the level of the expected discontinuity magnitude at the reference time, $\widehat{\mathcal{M}}_{\text{ref}}^{N,J}$, and the spot volatility, SV_s , described in Section 4. The factors are estimated from 3DTE, 4DTE, and 5DTE options, as the column groups indicate. Each column is a separate time-series regression across days, holding the forecast window fixed, with all variables scaled to unit standard deviation. The top panel uses the non-event days, the days excluding FOMC announcement and Minutes days and the COVID period (2020/03); the bottom panel uses the normal days, the non-event days whose daily realized variance falls between the 25th and 75th percentiles. The parentheses in the table report the Newey–West standard errors. The foot of each panel reports the R^2 (in %) and the number of days N . The sample period is from 2016/09 to 2023/12.

Non-Event									
Factor	3DTE			4DTE			5DTE		
	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$
α	0.42*** (0.08)	0.43*** (0.08)	0.50*** (0.07)	0.20 (0.18)	0.20 (0.17)	0.31** (0.15)	0.36*** (0.13)	0.37*** (0.13)	0.49*** (0.10)
$\Delta\widehat{\mathcal{M}}_s^{N,J}$	0.21*** (0.03)	0.25*** (0.04)	0.30*** (0.03)	0.17*** (0.02)	0.24*** (0.04)	0.24*** (0.03)	0.20*** (0.03)	0.28*** (0.05)	0.27*** (0.03)
$\widehat{\mathcal{M}}_{\text{ref}}^{N,J}$	0.13*** (0.05)	0.18*** (0.05)	0.23*** (0.04)	0.21** (0.09)	0.20*** (0.05)	0.23*** (0.04)	0.17** (0.07)	0.20*** (0.05)	0.25*** (0.05)
SV_s	0.60*** (0.04)	0.60*** (0.04)	0.68*** (0.05)	0.59*** (0.08)	0.62*** (0.08)	0.68*** (0.09)	0.61*** (0.06)	0.63*** (0.06)	0.67*** (0.06)
R^2	42.14	43.02	42.45	43.28	45.26	44.45	42.40	43.69	41.32
N	1,151	1,151	1,151	1,198	1,198	1,198	1,123	1,123	1,123
Normal									
Factor	3DTE			4DTE			5DTE		
	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$	$\Delta(s_1)$	$\Delta(s_2)$	$\Delta(s_3)$
α	1.07*** (0.14)	1.02*** (0.15)	1.11*** (0.16)	1.04*** (0.15)	0.93*** (0.16)	0.93*** (0.17)	1.10*** (0.14)	1.06*** (0.14)	1.13*** (0.16)
$\Delta\widehat{\mathcal{M}}_s^{N,J}$	0.21*** (0.06)	0.30*** (0.04)	0.27*** (0.05)	0.22*** (0.04)	0.25*** (0.04)	0.32*** (0.05)	0.21*** (0.06)	0.32*** (0.04)	0.29*** (0.05)
$\widehat{\mathcal{M}}_{\text{ref}}^{N,J}$	0.13*** (0.05)	0.18*** (0.06)	0.23*** (0.05)	0.14** (0.05)	0.19*** (0.06)	0.22*** (0.05)	0.10** (0.05)	0.17*** (0.05)	0.24*** (0.05)
SV_s	0.40*** (0.05)	0.42*** (0.05)	0.49*** (0.07)	0.37*** (0.05)	0.42*** (0.05)	0.51*** (0.07)	0.39*** (0.05)	0.42*** (0.05)	0.45*** (0.07)
R^2	19.46	19.12	20.49	17.96	19.29	21.78	16.64	16.32	16.93
N	585	585	585	615	615	615	555	555	555

★ denotes statistical significance at the 10% level, ★★ at the 5% level, and ★★★ at the 1% level.