

Decoupling Interday and Intraday Volatility Dynamics with Price Durations*

Yifan Li[†] Ingmar Nolte[‡] Sandra Nolte[§] Shifan Yu[¶]

This Version: May 5, 2025

Abstract

This paper introduces a novel framework for volatility estimation based on price durations with an adaptive price change threshold. This innovation allows us to disentangle daily and intraday volatility dynamics from price durations, which greatly simplifies the parametric modelling of price durations and hence leads to more accurate volatility estimators. Simulation results demonstrate superior finite-sample performance of our duration-based estimators for both spot and integrated volatility compared to some established methods. An empirical application based on intraday data for the SPDR S&P 500 ETF highlights the improved forecasting accuracy of our integrated volatility estimator within a standard daily volatility forecasting framework. Furthermore, an intraday analysis of spot volatility estimation shows that our method can capture the immediate and substantial impact of FOMC news announcements on market volatility.

Keywords: High-Frequency Data, Volatility Estimation, Price Durations, ACD Model

JEL Classifications: C14, C22, C41, C58

Mathematics Subject Classifications: 62F10, 62M10, 62P05

*We thank the Editor-in-Chief Robert Taylor, the Co-Editor Kim Christensen, and two anonymous referees for their constructive comments and suggestions. We also thank Manh Pham and Roberto Renò for their feedback and discussions. This research was supported by the Economic and Social Research Council (ESRC) North West Social Science Doctoral Training Partnership [Grant Number ES/P000665/1]. The usual disclaimer applies.

[†]Division of Accounting and Finance, Alliance Manchester Business School, University of Manchester, Booth Street West, Manchester M15 6PB, United Kingdom. Email: yifan.li@manchester.ac.uk.

[‡]Department of Accounting and Finance, Lancaster University Management School, Lancaster LA1 4YX, United Kingdom. Email: i.nolte@lancaster.ac.uk.

[§]Department of Accounting and Finance, Lancaster University Management School, Lancaster LA1 4YX, United Kingdom. Email: s.nolte@lancaster.ac.uk.

[¶]Corresponding author, Oxford-Man Institute of Quantitative Finance, University of Oxford, Eagle House, Walton Well Road, Oxford OX2 6ED, United Kingdom. Email: shifan.yu@eng.ox.ac.uk.

1 Introduction

Volatility is an important topic in financial econometrics and a crucial input for any asset pricing, portfolio allocation and risk management framework (Taylor, 2005). It is a latent process that describes the return variability over a local horizon which needs to be estimated from price observations. The increased availability of high-frequency financial data has motivated a shift of volatility estimation techniques from monthly or daily frequencies, such as the GARCH models (Engle, 1982; Bollerslev, 1986) and stochastic volatility models (Taylor, 1982, 1986, 1994), to various high-frequency volatility measures (Aït-Sahalia and Jacod, 2014). As the most representative and widely applied high-frequency volatility estimator, the realized volatility (RV) introduced by Andersen and Bollerslev (1998) is constructed by summing up all squared intraday log-returns, and is well-known to be a consistent and efficient estimator of the integrated variance (IV) of a univariate Itô semimartingale over fixed time intervals. The return-based RV estimator has well-established statistical properties and can be modified to accommodate more accurate volatility measures that are robust to various market frictions, e.g., Barndorff-Nielsen and Shephard (2004), Jacod et al. (2009), and Mancini (2009).

The seminal work of Engle and Russell (1998) provides a compelling alternative approach to return-based volatility estimation methods. Unlike the RV-type estimators that measure the magnitude of price changes over a given time interval, this alternative method measures the time it takes for the price to change by a certain size, i.e., a selected price change threshold. Motivated by Engle and Russell (1998), the duration-based volatility estimation has been further studied in Gerhard and Hautsch (2002), Andersen et al. (2008), Tse and Yang (2012), Fukasawa and Rosenbaum (2012), Vetter and Zwingmann (2017), Li et al. (2019, 2021), Hong et al. (2023), and Pelletier and Wei (2024), among others. Specifically, the parametric structure of the duration-based volatility estimators facilitates more flexible intraday inference for local volatility. As summarized by Tse and Yang (2012), the parametric duration-based estimators can benefit from the data beyond the estimation window to enhance parameter estimates, and potentially achieve more accurate volatility estimates relative to the conventional nonparametric return-based ones. Furthermore, the parametric structure facilitates the inclusion of other covariates such as seasonality and market microstructure covariates, which can not only improve the quality of volatility estimation, but also provide a framework to further explore the relation between volatility and other covariates at a high-frequency level.

The existence of market frictions requires the utilization of “not-too-finely” sampled data, which further restricts the data availability for both return- and duration-based methods in practice (Aït-Sahalia et al., 2005; Liu et al., 2015). For example, both Andersen et al. (2008) and Hong et al. (2023) recommend a moderate to large threshold to ensure a small number of durations relative to the available price observations on each day. Although the parametric structure offers the flexibility to estimate the duration models, e.g., the autoregressive conditional duration (ACD) model of Engle and Russell (1998), with the data beyond a specific day, the incorporation of intraday durations

across multiple days introduces complexities. Specifically, the durations obtained with the same threshold from different days will encompass different daily volatility dynamics, which leads to challenges in both model estimation and the analysis of volatility patterns. Although there has been extensive investigations on either daily or intraday volatility dynamics in the literature, a joint analysis of both is nearly infeasible due to their fundamentally different characteristics.

In this paper, we propose a new methodology to (i) disentangle the daily and intraday volatility dynamics inherent in the durations collected across multiple days, and (ii) parametrically estimate both spot and integrated volatility with price durations. Distinct from the existing literature which adopts a fixed threshold (see, e.g., [Tse and Yang, 2012](#); [Hong et al., 2023](#)), we employ a daily predicted threshold that adapts to changes in daily IV. As a result, interday volatility dynamics is subsumed into the daily adaptive thresholds, which homogenizes the price durations from different days by attenuating its long-run persistence, allowing more convenient parametric duration modelling that focuses on intraday volatility dynamics. Furthermore, we derive a relationship between spot volatility and the conditional density of durations for a semimartingale under some mild conditions, which formalizes the heuristic arguments in [Hautsch \(2011\)](#) and [Tse and Yang \(2012\)](#), and also extends the theoretical framework of [Pelletier and Wei \(2024\)](#).

Simulation results reveal that our new duration-based method delivers reliable finite-sample performance for both spot volatility and IV estimation. We benchmark our duration-based estimators against several localized return-based estimators in the spirit of [Foster and Nelson \(1996\)](#) for spot volatility estimation, and also compare them with various RV-type estimators for IV estimation. The results show that our duration-based estimators exhibit superior robustness to price jumps and variations in sampling frequencies, when compared to the selected competitors. In our first empirical application, we focus on the prediction of out-of-sample IV estimates of the SPDR S&P 500 ETF Trust (SPY) with the heterogeneous autoregressive (HAR) model of [Corsi \(2009\)](#). We find that HAR models based on our duration-based IV estimators outperform most of the selected benchmark models with significantly smaller forecast errors. Furthermore, we conduct an intraday analysis to evaluate the short-term impact of regular press releases by the Federal Reserve, i.e., the Federal Open Market Committee (FOMC) news announcements, on spot volatility. Our results confirm the well-documented fact that FOMC announcements at 14:00 have an instant and substantial impact on spot volatility (see, e.g., [Bollerslev et al., 2021, 2024](#)), which evidences the credibility of our duration-based spot volatility estimator.

The remainder of this paper is structured as follows: Section 2 details our estimation procedure, which includes both the daily threshold prediction and parametric volatility estimation. Section 3 presents an extensive Monte Carlo study to evaluate the finite-sample performance of our duration-based methods for spot and integrated volatility estimation. Section 4 provides empirical applications for both daily and intraday volatility of SPY. Section 5 concludes. Proofs and additional results can be found in the [Appendix](#).

2 Econometric Framework

We consider a one-dimensional underlying process $X = (X_t)_{t \geq 0}$ for the efficient logarithmic price of a financial asset. We assume that X follows a semimartingale defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$:

$$X_t = X_0 + \int_0^t \sigma_s dW_s + X'_t, \quad (1)$$

where t stands for time, $W = (W_t)_{t \geq 0}$ is a standard Brownian motion, $\sigma = (\sigma_t)_{t \geq 0}$ is a càdlàg \mathcal{F}_t -adapted process assumed to be locally bounded and bounded away from zero. We assume that X is observed on $[0, T) \cup [T, 2T) \cup \dots \cup [(d-1)T, dT) \cup \dots$, where the interval $[(d-1)T, dT)$ represents the d -th trading day of length $T > 0$. Price movements during market closures are modelled as jumps occurring at dT , i.e., $\Delta X_d = X_{dT} - X_{dT-}$, and are cumulatively incorporated into the discontinuous component $X'_t = \sum_{d=1}^{\lfloor t/T \rfloor} \Delta X_d$. We assume $T = 1$ for ease of notation in subsequent discussions.

Remark 1. We do not consider price jumps during the regular trading sessions due to the natural robustness of duration-based methods to finite-activity jumps. Some relevant discussions can be found in Andersen et al. (2008), Tse and Yang (2012), Hong et al. (2023), Pelletier and Wei (2024), and Li et al. (2025). Empirical evidence from the analysis of some liquid S&P 500 stocks in Christensen et al. (2025) reveals a significantly higher occurrence of price jumps during after-hours sessions, which often coincide with the release of earnings announcements.

We are interested in the estimation of spot variance σ_t^2 for some t as well as the integrated variance (IV) over some interval $[s, t]$:

$$V(s, t) = \int_s^t \sigma_u^2 du. \quad (2)$$

When the interval $[s, t] = [d-1, d]$ is one trading day, we write $V_d = V(d-1, d)$ as the IV for day d and $V(t) = V(0, t)$ to denote the IV process up to time t . To construct an estimator for $V(t)$, we extend the duration-based methods of Tse and Yang (2012) and Hong et al. (2023). In contrast to those return-based estimators that fix a time interval Δ and measure the change in X , the duration-based estimators fix a threshold $\delta > 0$ and measure the durations when X increases or decreases by δ . Specifically, we generalize the price duration sampling (PDS) method of Hong et al. (2023) by incorporating a time-varying stopping rule. We choose a sequence of “daily” thresholds $(\delta_d)_{d=1,2,\dots}$, where each δ_d is positive and adapted to the information \mathcal{F}_{d-1} available up to time $d-1$. We then sample price observations within each trading day based on the following stopping rule:

$$\tau_{d,0} = d-1, \quad \tau_{d,i} = \inf_{\tau_{d,i-1} < t \leq d} \{|X_t - X_{\tau_{d,i-1}}| \geq \delta_d\}, \quad (3)$$

with the convention that $\inf\{\emptyset\} = \infty$. Consequently, for each day d , we obtain a sequence of sampling times $(\tau_{d,i})_{0 \leq i \leq N_d}$, where $\tau_{d,i}$ denotes the i -th price event, $x_{d,i} = \tau_{d,i} - \tau_{d,i-1}$ the i -th inter-event price duration, and $N_d = \sum_{i \geq 0} \mathbb{1}_{\{\tau_{d,i} \in (d-1, d]\}}$ counts the number of price events within

the d -th day.

In essence, when the realized sample path $X(\omega)$ is fully observable, the absolute return between each pair of consecutive sampled price observations equals δ_d , i.e., $|X_{\tau_{d,i}} - X_{\tau_{d,i-1}}| = \delta_d$ for all $1 \leq i \leq N_d$. Consequently, N_d can be interpreted as the frequency of price path changes by δ_d , which is proportional to the nonparametric duration-based IV estimator proposed by [Hong et al. \(2023\)](#). This argument is formalized by Theorem 1 of [Hong et al. \(2023\)](#), which implies that

$$N_d = \frac{V_d}{\delta_d^2} + M_d + o_p(1), \quad (4)$$

where $(M_d)_{d=1,2,\dots}$ is a Gaussian martingale difference sequence, and $o_p(1)$ vanishes as $\delta_d \rightarrow 0$. By taking conditional expectations on both sides of Eq. (4), the \mathcal{F}_{d-1} -adaptedness of δ_d further implies:

$$\mathbb{E}[N_d | \mathcal{F}_{d-1}] = \frac{\mathbb{E}[V_d | \mathcal{F}_{d-1}]}{\delta_d^2} + o(1). \quad (5)$$

Therefore, for a fixed δ , $\mathbb{E}[N_d | \mathcal{F}_{d-1}] \propto \mathbb{E}[V_d | \mathcal{F}_{d-1}]$, implying that daily price event counts effectively capture the volatility dynamics on a daily horizon. This can be further modelled parametrically with the durations or intensities of point processes ([Engle and Russell, 1998](#); [Tse and Yang, 2012](#); [Hong et al., 2023](#)).

As an important innovation of this paper, we notice that one does not need to choose a fixed δ . Instead, we introduce an adaptive choice of $\delta_d = \sqrt{K^{-1} \mathbb{E}[V_d | \mathcal{F}_{d-1}]}$ for some constant K , which ensures that the expected number of price events on day d is approximately K , i.e., $\mathbb{E}[N_d | \mathcal{F}_{d-1}] \approx K$.

The adaptive choice of δ_d has two advantages over a constant δ . Firstly, it offers a natural mechanism for controlling daily sampling frequencies. In practice, the full trajectory of $X(\omega)$ is not observable, and is instead recorded on a grid of discrete times with the contamination of market microstructure noise. Previous studies on duration-based volatility estimation recommend selecting a moderate to large δ to ensure that the sampling frequency is sufficiently low relative to the total number of observations ([Andersen et al., 2008](#); [Li et al., 2021](#); [Hong et al., 2023](#)). However, achieving this “not-too-finely” sampling is feasible only on average with a fixed δ . As shown in Eq. (5), the expected daily sampling frequency depends on the daily IV, which can vary substantially across multiple days. With an appropriate choice of K (for example, $K = 78$ represents an average price duration of five minutes), the adaptive threshold δ_d allows for direct control over the expected sampling frequency for each day, which provides adaptive protection against market imperfections on a daily basis.

Secondly, the construction of δ_d implies that $K\delta_d^2$ can be interpreted as the optimal forecast of V_d in terms of mean squared error (MSE). The persistence of volatility dynamics at the daily level has been extensively studied in the literature ([Corsi, 2009](#); [Gatheral et al., 2018](#)), and these benchmark models provide empirically reliable methods to construct δ_d that fully reflects the daily volatility dynamics. For example, since the squared threshold is essentially a scaled IV proxy, some predictive models for one-day-ahead RV, such as the HAR model of [Corsi \(2009\)](#), can be directly applied. As

a result, the sampling times $(\tau_{d,i})_{0 \leq i \leq N_d}$ within each day only preserve intraday volatility dynamics, as the interday persistence in the IV process is subsumed into the adaptive thresholds $(\delta_d)_{d=1,2,\dots}$. This decomposition allows us to model interday and intraday volatility separately, which greatly increases the flexibility of point-process-based parametric volatility models, such as those proposed by [Hong et al. \(2023\)](#) and [Pelletier and Wei \(2024\)](#). It also considerably simplifies the econometric analysis, offering a clearer and more intuitive understanding of volatility dynamics over extended periods spanning multiple days.

2.1 Spot and Integrated Volatility Estimation with Price Durations

We proceed to explain how we estimate both spot and integrated variances from the observed price durations $(x_{d,i})_{1 \leq i \leq N_d}$ based on some \mathcal{F}_{d-1} -adapted threshold. Let $\mathcal{F}_{d,i}$ denote the filtration generated by X up to the sampling time $\tau_{d,i}$. We will demonstrate that the $\mathcal{F}_{d,i-1}$ -conditional density of $x_{d,i}$ is intrinsically linked to the spot variance σ^2 of the continuous martingale X .

To this end, we introduce some additional notation: Let $V_{d,i} = V(d-1, d-1 + \tau_{d,i})$ denote the IV accumulated up to the i -th price event on day d , and thus $\Delta_{d,i}V = V_{d,i} - V_{d,i-1}$ represents the i -th duration in the IV clock. The sequence of IV increments $(\Delta_{d,i}V)_{1 \leq i \leq N_d}$ has the following important property:

Proposition 1. Assume $(\Delta_{d,i}V)_{1 \leq i \leq N_d}$ is generated from the price model in Eq. (1) with the \mathcal{F}_{d-1} -adapted threshold δ_d . Then it holds that

$$\Delta_{d,i}V = \delta_d^2 Z_i, \quad (6)$$

where $(Z_i)_{1 \leq i \leq N_d}$ is a sequence of independent and identically distributed (i.i.d.) positive random variables such that for all i , $Z_i \stackrel{d}{=} \inf_{t>0} \{|B_t| \geq 1\}$ for some standard Brownian motion B with Z_i independent of $\mathcal{F}_{d,i-1}$.

Remark 2. Proposition 1 stems from the well-known Dambis-Dubins-Schwarz theorem that all continuous martingales are time-changed Brownian motions under the IV clock, or the business time ([Barndorff-Nielsen and Shiryaev, 2015](#)). As the price events commute with time changes, $(\Delta_{d,i}V)_{1 \leq i \leq N_d}$ is, up to a constant scaling, identical in distribution to $(Z_i)_{1 \leq i \leq N_d}$. The i.i.d.-ness thus follows from the strong Markov property and time homogeneity of the Brownian motion. The density of Z_i is well-known in the literature with the following probability density function (PDF) and cumulative distribution function (CDF):

$$f_Z(z) = \sum_{k=-\infty}^{\infty} \frac{2(1+4k)}{\sqrt{2\pi}z^{3/2}} e^{-\frac{(1+4k)^2}{2z}}, \quad F_Z(z) = 2 - 2 \sum_{k=-\infty}^{\infty} \operatorname{erf} \frac{1+4k}{\sqrt{2z}}, \quad (7)$$

where $\operatorname{erf} x = (2/\sqrt{\pi}) \int_0^x e^{-t^2} dt$ is the error function ([Andersen et al., 2008](#)). In particular, we have $\mathbb{E}[Z_i] = \mathbb{E}[Z_i | \mathcal{F}_{d,i-1}] = 1$, which implies that the expected IV increment between two price events is precisely δ_d^2 for all continuous martingales, i.e., $\mathbb{E}[\Delta_{d,i}V | \mathcal{F}_{d,i-1}] = \delta_d^2$.

A key insight from Proposition 1 is that the IV-based time change converts $x_{d,i}$ into $\Delta_{d,i}V$ for each i , where the latter is, conditional on $\mathcal{F}_{d,i-1}$, an i.i.d. duration in the IV clock. Consequently, it is natural to expect that the $\mathcal{F}_{d,i-1}$ -conditional distribution of $x_{d,i}$ contains information about IV. To formalize this link, we define a counting process $N_d(t) = \sum_{i \geq 1} \mathbb{1}_{\{\tau_{d,i} \leq t\}}$ which counts the number of price events on day d up to time $t \in [d-1, d]$, and a piecewise constant price process $\bar{X}_t = X_{\tau_{d,N_d(t)}}$, which is constant on each inter-event interval $[\tau_{d,i-1}, \tau_{d,i})$. Let $\bar{\mathcal{F}}_t \subset \mathcal{F}_t$ denote the natural filtration generated by \bar{X}_t up to time t . Similarly, let $\bar{\mathcal{F}}_{d,i} = \bar{\mathcal{F}}_{\tau_{d,i}}$ denote the filtration at the i -th sampling time. With these definitions in place, we establish the following results:

Proposition 2. Assume the same conditions in Proposition 1 hold and further assume that $V(t)$ is adapted to $\bar{\mathcal{F}}_t$. Let $\Delta_{d,i}V(h) = V(\tau_{d,i-1}, \tau_{d,i-1} + h)$ denote the IV increment over the interval $[\tau_{d,i-1}, \tau_{d,i-1} + h]$. Let $f(h|\bar{\mathcal{F}}_{d,i-1})$ and $F(h|\bar{\mathcal{F}}_{d,i-1})$ denote the $\bar{\mathcal{F}}_{d,i-1}$ -conditional PDF and CDF of $x_{d,i}$, respectively. For all d, i , it holds for all $h \in (0, x_{d,i}]$ that:

$$\Delta_{d,i}V(h) = \delta_d^2 G_{d,i}(h), \quad (8)$$

where $G_{d,i}(h) = F_Z^{-1}(F(h|\bar{\mathcal{F}}_{d,i-1}))$ and $F_Z^{-1}(\cdot)$ is the inverse function of the CDF in Eq. (7). Furthermore, for almost all $h \in (0, x_{d,i}]$, we have:

$$\sigma_{\tau_{d,i-1}+h}^2 = \frac{\delta_d^2 f(h|\bar{\mathcal{F}}_{d,i-1})}{f_Z(G_{d,i}(h))}. \quad (9)$$

Remark 3. The use of the restricted filtration $\bar{\mathcal{F}}_t$ instead of \mathcal{F}_t as well as the $\bar{\mathcal{F}}_t$ -adaptedness of $V(t)$ require some elaboration. First, to reflect the PDS procedure on X , it is natural to consider the filtration generated only by sampled observations. This restriction is implicit in the existing work (Tse and Yang, 2012; Hong et al., 2023), which is intended to avoid using the full (potentially noisy) price paths in \mathcal{F}_t . The adaptedness of $V(t)$ to $\bar{\mathcal{F}}_t$ also ensures the validity of Proposition 1.¹ More importantly, it ensures that $V(t)$ on $[\tau_{d,i-1}, \tau_{d,i})$ is $\bar{\mathcal{F}}_{d,i-1}$ -adapted. This allows us to construct an observation-driven model for a continuous-time stochastic process, which can be easily estimated based on standard econometric tools. We show in our simulation section that a $\bar{\mathcal{F}}_t$ -adapted volatility process can still provide accurate approximation to $V(t)$ fully driven by a stochastic volatility process that violates the $\bar{\mathcal{F}}_t$ -adaptedness restriction.

Proposition 2 leads to the following representation of IV over any finite interval $[s, t]$ in day d :

$$V(s, t) = \delta_d^2 \left[\left(\sum_{i=N_d(s)}^{N_d(t)} G_{d,i}(x_{d,i}) \right) - G_{d,N_d(s)}(s - \tau_{d,N_d(s)}) + G_{d,N_d(t)+1}(t - \tau_{d,N_d(t)}) \right], \quad (10)$$

where the last two terms correct for the (left and right) edge effects due to a mismatch between the interval $[s, t]$ and $[\tau_{d,N_d(s)}, \tau_{d,N_d(t)}]$. Taking $[s, t] = [d-1, d]$ gives the IV for day d , where the left correction is no longer needed. The right correction, termed the end-of-day correction in Hong et al. (2023), typically exhibits a smaller magnitude compared to the leading term and is often ignored in

practical applications.

It is worth noting that Proposition 2 only identifies the spot variance σ_t^2 for almost all t up to a Lebesgue null set. Intuitively, this is due to the fact that the spot variance process is identified through IV, whose value does not change by altering σ_t^2 on a Lebesgue null set. This also indicates that the point-wise result in Proposition 2 may not be very informative about the spot volatility at finitely many t . Instead, we propose to approximate σ_t by the localized IV over the interval $[t - h, t]$ for some small $h > 0$:

$$\sigma_t = \sqrt{\frac{1}{h}V(t - h, t)}, \quad (11)$$

where $V(s, t)$ is given in Eq. (10). This smooths out any potential point-wise divergence from Eq. (9), which is adopted in our simulation and empirical analyses.

As an interesting special case of Proposition 2, we derive a condition for σ^2 to be almost everywhere piecewise constant:

Corollary 1. Suppose $(\gamma_{d,i})$ is a sequence of $\overline{\mathcal{F}}_{d,i-1}$ -adapted positive random variables, such that

$$x_{d,i} = \frac{\delta_d^2}{\gamma_{d,i}} Z_i, \quad (12)$$

for all d and i , then it holds for almost all $h \in (0, x_{d,i}]$ that

$$\Delta_{d,i}V(h) = h\gamma_{d,i} \quad \text{and} \quad \sigma_{\tau_{d,i-1}+h}^2 = \gamma_{d,i}. \quad (13)$$

Corollary 1 implies that σ^2 is almost everywhere piecewise constant on $[\tau_{d,i-1}, \tau_{d,i}]$ under the following two conditions: (i) The duration $x_{d,i}$ satisfies Eq. (12) for all d and i , and (ii) $\gamma_{d,i}$ is $\overline{\mathcal{F}}_{d,i-1}$ -adapted, which means $x_{d,i}$ is proportional to Z_i conditional on $\overline{\mathcal{F}}_{d,i-1}$. This is at odds with some of the existing point-process-based volatility estimators in the literature, to which we shall turn. The existing methods, e.g., Hautsch (2011), Tse and Yang (2012), and Hong et al. (2023), typically adhere to two equivalent methodologies: the duration-based method of Engle and Russell (1998), and the intensity-based method of Gerhard and Hautsch (2002). Both methods employ a heuristic argument that each price event contributes δ_d^2 to V_d , and thus they decompose the spot variance multiplicatively as the product of the squared threshold δ_d^2 and the intensity or hazard rate of price events, i.e.,

$$\sigma_{\tau_{d,i-1}+h}^2 = \delta_d^2 \lambda(h|\overline{\mathcal{F}}_{d,i-1}), \quad h \in [0, x_{d,i}), \quad (14)$$

where $\lambda(h|\overline{\mathcal{F}}_{d,i-1})$ is the $\overline{\mathcal{F}}_{d,i-1}$ -conditional hazard function of the price events at time $\tau_{d,i-1} + h$, as defined in Daley and Vere-Jones (2003):

$$\lambda(h|\overline{\mathcal{F}}_{d,i-1}) = \frac{f(h|\overline{\mathcal{F}}_{d,i-1})}{1 - F(h|\overline{\mathcal{F}}_{d,i-1})}. \quad (15)$$

This model offers a convenient formulation for the spot variance in a multiplicative structure.

Despite this relatively simpler specification of σ_t^2 , this decomposition cannot hold for the

continuous martingale model in Eq. (1).² With the spot variance in Eq. (14), the IV increment between $\tau_{d,i-1}$ and $\tau_{d,i}$ is given by

$$\Delta_{d,i}V = \int_{\tau_{d,i-1}}^{\tau_{d,i}} \sigma_s^2 ds = \delta_d^2 \int_0^{x_{d,i}} \lambda(h|\mathcal{F}_{d,i-1}) dh = -\delta_d^2 \ln(1 - F(x_{d,i}|\bar{\mathcal{F}}_{d,i-1})). \quad (16)$$

In this case, with a copula transformation argument, $\Delta_{d,i}V$ follows an exponential distribution with the intensity parameter δ_d^{-2} , which clearly contradicts Proposition 1. Specifically, it implies that the likelihood of a (time-changed) Brownian motion exiting a symmetric barrier $[-\delta_d, \delta_d]$ is time-invariant. Such a memoryless property is impossible due to the continuity of Brownian diffusions, see, e.g., Eq. (3.0.2) in Borodin and Salminen (2002). Therefore, $\tilde{\sigma}^2$ cannot coincide with the spot variance process of a continuous martingale.

Taking a different approach, Pelletier and Wei (2024) adopt the local volatility approximation of Andersen et al. (2008) and assume that σ is piecewise constant on all intervals of the form $(\tau_{d,i-1}, \tau_{d,i})$. Jointly with Proposition 1, it implies that, for each $1 \leq i \leq N_d$,

$$x_{d,i} \stackrel{d}{=} \frac{\delta_d^2}{\sigma_{\tau_{d,i-1}}^2} Z_i. \quad (17)$$

The difference between the above result and our Corollary 1 is two-fold. On the one hand, Pelletier and Wei (2024) allow $\sigma_{\tau_{d,i-1}}^2$ to depend on concurrent information up to time $\tau_{d,i}$ through an additional stochastic component, while our Corollary 1 requires $\sigma_{\tau_{d,i-1}}^2$ to be adapted to $\bar{\mathcal{F}}_{d,i-1}$. This can be considered as a stochastic extension of our approach, which can potentially provide more flexibility to model the volatility dynamics. However, the resulting model becomes fully parameter-driven, which requires computationally intensive estimation techniques and considerably complicates its empirical implementation (Koopman et al., 2016). On the other hand, Proposition 2 demonstrates that the piecewise constant assumption can be relaxed to a piecewise adaptedness condition, which allows for a more flexible specification for σ even within our observation-driven framework. We will thus focus on Proposition 2 in developing volatility models and leave the generalization to a fully parameter-driven model for future research.

2.2 A Joint Model for Daily and Intraday Volatility Dynamics

In this section, we describe our econometric model for both the daily thresholds $(\delta_d)_{1 \leq d \leq D}$ and the intraday price durations $x_{d,i}$, where D denotes the total number of trading days in the in-sample period. We start with a model for $(\delta_d)_{1 \leq d \leq D}$: Let $K \in \mathbb{Z}^+$ denote a predetermined “expected” sampling frequency. The realized K -adaptive threshold on the d -th day is defined as

$$\delta_d = \sup_{\delta > 0} \operatorname{argmin}_{\delta} |K - N_d|, \quad (18)$$

where the supremum is taken to ensure the uniqueness of δ_d . Intuitively, δ_d is the largest threshold that generates a sampling frequency N_d closest to K , which can be easily constructed from the

price observations on day d in practice. Note that $K\delta_d^2$ actually serves as a daily IV estimator and can be treated as a proxy for V_d . Based on a burn-in sample of $(\delta_d)_{-h \leq d \leq 0}$, we estimate the HAR model of Corsi (2009):

$$\delta_d^2 = \omega_0 + \omega_1 \delta_{d-1}^2 + \omega_2 \sum_{i=1}^5 \delta_{d-i}^2 + \omega_3 \sum_{i=1}^{22} \delta_{d-i}^2 + \epsilon_d, \quad (19)$$

which can be easily estimated via ordinary least squares (OLS). The HAR model is designed to parsimoniously capture the dependence structure of IV across different horizons, and therefore aims to approximate its long memory which has been extensively confirmed by the empirical literature. Renowned for its consistent and remarkable predictive performance, the HAR model serves as the predominant benchmark in modelling and forecasting daily IV dynamics. With the parameter estimates $\hat{\omega} = (\hat{\omega}_0, \hat{\omega}_1, \hat{\omega}_2, \hat{\omega}_3)$, we set $\hat{\delta}_1^2$ as the one-step-ahead forecast of δ_1^2 with the optimal MSE:

$$\hat{\delta}_1^2 = \mathbb{E}[\delta_1^2 | \mathcal{F}_0] = \hat{\omega}_0 + \hat{\omega}_1 \delta_0^2 + \hat{\omega}_2 \sum_{i=1}^5 \delta_{1-i}^2 + \hat{\omega}_3 \sum_{i=1}^{22} \delta_{1-i}^2, \quad (20)$$

and the values of $(\hat{\delta}_d)_{2 \leq d \leq D}$ can be obtained with recursive model estimation and prediction in a rolling-window fashion. The construction of $\hat{\delta}_d$ ensures that $\hat{\delta}_d^2 = \mathbb{E}[\delta_d^2 | \mathcal{F}_{d-1}] \approx K^{-1} \mathbb{E}[V_d | \mathcal{F}_{d-1}]$, which is the desired threshold adaptive to daily volatility dynamics.

With the sequence of thresholds $(\hat{\delta}_d)_{1 \leq d \leq D}$, we obtain all durations across all D days, which can be modelled parametrically with standard duration-based point process models (Hautsch, 2011). For example, we consider the log-ACD-GARCH model of Allen et al. (2008) as follows:

$$\begin{aligned} \ln x_{d,i} &= \Psi_{d,i} + s_{d,i} + \varepsilon_{d,i}, \quad \varepsilon_{d,i} = \sqrt{h_{d,i}} u_{d,i}, \\ \Psi_{d,i} &= \sum_{j=1}^p \phi_j \Psi_{d,i-j} + \sum_{j=1}^q \theta_j \varepsilon_{d,i-j}, \\ h_{d,i} &= \tilde{s}_{d,i} + \sum_{j=1}^{p^*} \alpha_j \varepsilon_{d,i-j}^2 + \sum_{j=1}^{q^*} \beta_j h_{d,i-j}, \end{aligned} \quad (21)$$

where $(u_{d,i})$ is a sequence of i.i.d. random variables with zero mean and unit variance, with a parametric PDF $f_u(x; \gamma)$ governed by the parameter vector γ . The processes $(\Psi_{d,i})$ and $(h_{d,i})$ are standard conditional mean and variance specifications used in ACD- and GARCH-type models, which are both $\bar{\mathcal{F}}_{d,i-1}$ -adapted. The variables $s_{d,i}$ and $\tilde{s}_{d,i}$ are also $\bar{\mathcal{F}}_{d,i-1}$ -adapted components that capture the seasonality in the mean and variance of log-durations, respectively, and are specified in flexible Fourier forms following Andersen and Bollerslev (1997). For example, a Q -th-order flexible-Fourier-form specification for $s_{d,i}$ is given by

$$s_{d,i} = \nu_0 + \nu_1 \bar{\tau}_{d,i-1} + \nu_2 \bar{\tau}_{d,i-1}^2 + \sum_{j=1}^Q (\nu_{c,j} \cos(2\pi j \cdot \bar{\tau}_{d,i-1}) + \nu_{s,j} \sin(2\pi j \cdot \bar{\tau}_{d,i-1})), \quad (22)$$

where the parameters $\{\nu_0, \nu_1, \nu_2, \nu_{c,1}, \dots, \nu_{c,Q}, \nu_{s,1}, \dots, \nu_{s,Q}\}$ for $s_{d,i}$, and similarly $\{\tilde{\nu}_0, \tilde{\nu}_1, \tilde{\nu}_2, \tilde{\nu}_{c,1}, \dots, \tilde{\nu}_{c,Q}, \tilde{\nu}_{s,1}, \dots, \tilde{\nu}_{s,Q}\}$ for $\tilde{s}_{d,i}$, are jointly estimated with other model parameters. We follow Hautsch (2011) to reset the autoregressive structure of $\ln x_{d,i}$ and $h_{d,i}$ in Eq. (21) at the beginning of each trading day, as we do not expect the end-of-day duration to impact the first duration of the subsequent day.

Given that the log-ACD-GARCH- s, \tilde{s} model in Eq. (21) is in essence a variant of the celebrated GARCH model of Engle (1982) and Bollerslev (1986), its related estimation techniques and asymptotic properties are well-understood in the literature. Specifically, since the $\mathcal{F}_{d,i-1}$ -conditional density of $x_{d,i}$ is determined by $f_u(x; \gamma)$, all model parameters can be jointly estimated with standard maximum likelihood estimation (MLE). Under the correct model specification and standard regularity conditions (see, e.g., Amemiya, 1985), the MLE estimator is consistent and asymptotically normal as the number of days $D \rightarrow \infty$.

Furthermore, the volatility dynamics implied by the log-ACD-GARCH- s, \tilde{s} model in Eq. (21) need some discussions. The $\overline{\mathcal{F}}_{d,i-1}$ -adaptedness of $\Psi_{d,i}$, $s_{d,i}$, and $h_{d,i}$ indicates that

$$\mathbb{E}[x_{d,i} | \overline{\mathcal{F}}_{d,i-1}] = e^{\Psi_{d,i} + s_{d,i}} \mathbb{E}[e^{\sqrt{h_{d,i}} u_{d,i}} | \overline{\mathcal{F}}_{d,i-1}]. \quad (23)$$

The term $\exp(\Psi_{d,i} + s_{d,i})$ multiplicatively captures the autoregressive structure and seasonality of intraday price durations. The GARCH-type conditional variance $h_{d,i}$ further allows non-multiplicative autoregressive structure and seasonality to be modelled through the conditional moment generating function of $u_{d,i}$. As a special case under the conditions in Corollary 1, if $u_{d,i} \sim \ln Z_i - \mathbb{E}[\ln Z_i]$, where Z_i is defined in Proposition 1, and $h_{d,i} = 1$, then by Eq. (17) we have

$$x_{d,i} = e^{\Psi_{d,i} + s_{d,i}} Z_i \quad \text{and} \quad \sigma_{\tau_{d,i-1}}^2 = \delta_d^2 e^{-\Psi_{d,i} - s_{d,i}}, \quad (24)$$

where $\ln Z_i$ can be subsumed into the seasonality factor $s_{d,i}$. This implies that σ is piecewise constant on $[\tau_{d,i-1}, \tau_{d,i})$, and can be viewed as a càdlàg (i.e., right-continuous with left limits) process in continuous time that is only updated when the sampling occurs. The changed value of σ on $\tau_{d,i}$ depends on both the self-dependence structure $\exp(-\Psi_{d,i})$ and the seasonal pattern $\exp(-s_{d,i})$.

By allowing for different distributional assumptions of $u_{d,i}$ with the further inclusion of the conditional variance component $h_{d,i}$, the model permits time-varying spot volatility both between and within price durations, which enriches the intraday volatility dynamics implied by duration-based models. For example, Fig. 1 illustrate the impact of different distribution assumptions of $u_{d,i}$ to the shape of spot volatility and IV as time elapses from the previous price event. In particular, the dark dash lines in Fig. 1 corroborate the result in Corollary 1 that when $x_{d,i} \stackrel{d}{=} Z_i$, the spot volatility is constant between price durations, implying a linearly increasing IV as a function of time. By changing the distributional assumption of $x_{d,i}$, our method allows different shapes of the “baseline” spot volatility that evolves solely with respect to the time elapse since the last event, which echoes the $\overline{\mathcal{F}}_t$ -adaptedness property of the spot volatility. Therefore, the choice of the density of $x_{d,i}$ mirrors the choice of the baseline function in the autoregressive conditional intensity (ACI)

model (Russell, 1999; Hautsch, 2011), which specifies the evolution of spot volatility between price events as a deterministic function of time.

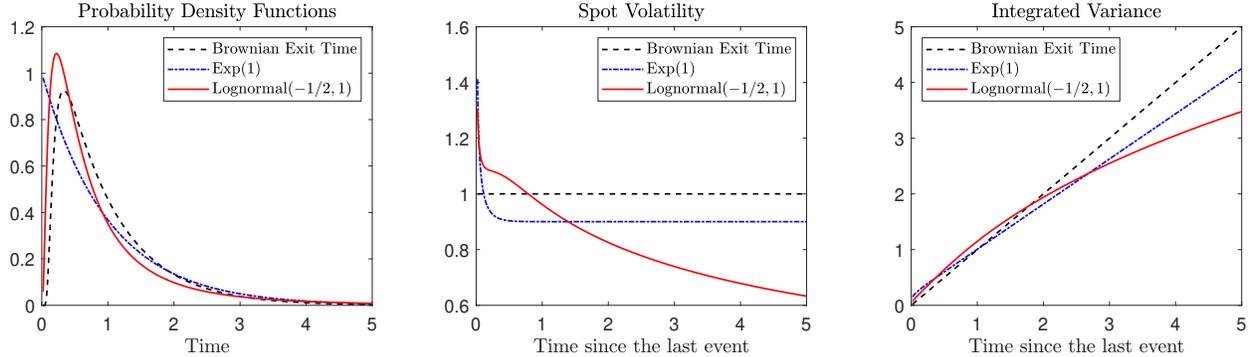


Figure 1: Examples of spot volatilities and IV increments over $[\tau_{d,i-1}, \tau_{d,i-1} + h]$ with $h \in (0, x_{d,i}]$, assuming a unit sampling threshold and three different distributional assumptions for durations $x_{d,i}$ normalized to have a unit mean: (i) the Brownian exit time Z_i defined in Proposition 1, (ii) the unit exponential distribution $\text{Exp}(1)$, and (iii) the $\text{Lognormal}(-1/2, 1)$ distribution.

3 Simulation Results

This section contains a Monte Carlo study to examine the finite-sample performance of our duration-based volatility estimators, which corresponds to the results developed in Sections 2.1 and 2.2.

3.1 Simulation Design

We simulate a Heston model for the efficient price process X (Heston, 1993):

$$\begin{aligned} dX_t &= \mu dt + \sigma_t dW_{1,t} + dJ_t, \quad \sigma_t = \check{\sigma}_t \gamma_t, \\ d\check{\sigma}_t^2 &= \alpha (\theta_d - \check{\sigma}_t^2) dt + \eta \check{\sigma}_t dW_{2,t}, \end{aligned} \tag{25}$$

where W_1 and W_2 are standard Brownian motions with $\text{Corr}(W_{1,t}, W_{2,t}) = \rho$, and J is a compound Poisson process, i.e.,

$$J_t = \sum_{i=1}^{N_t} D_i, \tag{26}$$

where N is a Poisson process with rate λ , and D_i follows a normal distribution $\mathcal{N}(0, \zeta^2)$. For the spot volatility $\sigma_t = \check{\sigma}_t \gamma_t$, we follow Hasbrouck (1999), Andersen et al. (2012), and Christensen et al. (2018) to model the diurnal pattern of intraday volatility in γ_t with a sum of two exponentials:

$$\gamma_t = C + Ae^{-a_1 t} + Be^{-a_2(1-t)}. \tag{27}$$

We set $A = 0.75$, $B = 0.25$, $C = 0.88929188$, and $a_1 = a_2 = 10$. This realistically calibrated specification produces a pronounced, asymmetric reverse J-shape in $\sigma_{u,t}$, with variance at $t = 0$

(resp. $t = 1$) more than three times (resp. about 1.5 times) the midday variance ($t = 1/2$).

The annualized parameters for Eq. (25) are fixed at $(\mu, \alpha, \theta_0, \eta, \rho) = (0.05, 5, 0.16, 0.5, -0.5)$, where the volatility parameters satisfy the Feller's condition $2\alpha\theta_0 \geq \eta^2$ which ensures the positivity of σ . The parameter choices follow both Aït-Sahalia and Jacod (2009) and Aït-Sahalia et al. (2012), which are calibrated according to the empirical estimates in Aït-Sahalia and Kimmel (2007). Specifically, for the annualized daily variance parameter θ_d , we assume a HAR structure as follows:

$$\theta_d = \omega_0 + \omega_1\theta_{d-1} + \omega_2 \sum_{i=1}^5 \theta_{d-i} + \omega_3 \sum_{i=1}^{22} \theta_{d-i} + \epsilon_d^\theta, \quad \text{with } \epsilon_d^\theta \sim \text{i.i.d. } \mathcal{N}(0, \kappa^2), \quad (28)$$

where $(\omega_0, \omega_1, \omega_2, \omega_3, \kappa) = (0.03, 0.2, 0.4/5, 0.2/22, 0.03)$. Additionally, the process J simulated with $\lambda = 1/5$ and $\varsigma = 2\%$ implies an average of one jump per week, and the jump variation is about 20% of the daily IV on average, which is consistent with Andersen et al. (2023). Fig. 2 illustrates the intraday variation of returns and annualized RVs of a simulated path in each one-minute interval. The return variation exhibits an asymmetric U-shaped or reverse J-shaped pattern over the trading hours, which is in line with some prior empirical findings (Harris, 1986; Wood et al., 1985; Andersen and Bollerslev, 1997; Christensen et al., 2018; Andersen et al., 2018, 2019, 2024).

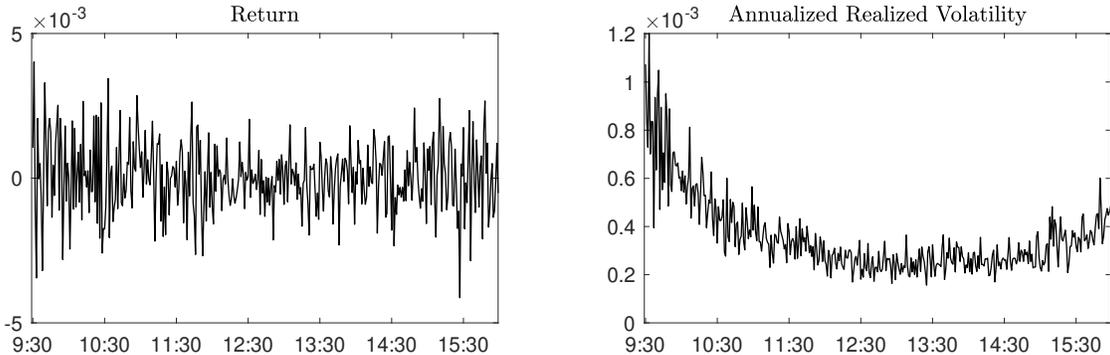


Figure 2: Intraday variation of returns and annualized RVs of a simulated Heston process. The tick-level observations are simulated with the Heston model in Eq. (25), with a pronounced, asymmetric reverse J-shape exhibited in spot volatility. The returns and annualized RVs are computed at a granularity of one minute.

We simulate second-by-second observations for 500 consecutive days. The realized K -adaptive thresholds over the first 200 days, i.e., the burn-in sample $\{-h, \dots, 0\}$, are used to estimate the HAR model in Eq. (19) and predict the threshold $\hat{\delta}_1$ for the 201-st day. In this section, we focus on both the spot volatility and IV estimation for the last 300 days.

3.2 Model Estimation

For the K -adaptive thresholds, we consider three different values of K , i.e., $K = 78, 39$, and 26 , for the burn-in sample, which correspond to similar levels of sparsity of 5, 10 and 15-minute calendar-time sampling (CTS), respectively. The durations observed in the last 300 days are obtained with the predicted thresholds from the HAR model in Eq. (19) estimated in a rolling-window fashion.

As mentioned in Section 2, by controlling the daily K -adaptive thresholds, we remove the interday volatility dynamics from the sequence of durations across all 300 days. To obtain the conditional CDF of durations, we estimate the log-ACD-GARCH- s, \tilde{s} model in Eq. (21) with MLE (with $u_{d,i}$ assumed as a standard Gaussian white noise). Table 1 reports the mean and standard deviation of the estimated parameters for the log-ACD-GARCH- s, \tilde{s} model, where we select the lags (1, 1) for both the ACD and GARCH parts, with second-order flexible-Fourier-form specifications for both seasonality terms s and \tilde{s} .³

Table 1: Parameter estimates of log-ACD-GARCH- s, \tilde{s} model

Parameters	$K = 78$	$K = 39$	$K = 26$	Parameters	$K = 78$	$K = 39$	$K = 26$
ϕ_1	0.8776 (0.0134)	0.7291 (0.0292)	0.7283 (0.0314)	α_1	0.0012 (0.0044)	0.0029 (0.0059)	0.0017 (0.0081)
θ_1	-0.8018 (0.0185)	-0.6315 (0.0311)	-0.6357 (0.0345)	β_1	0.8418 (0.2040)	0.8739 (0.1270)	0.8009 (0.0692)
ν_0	0.2970 (0.0616)	0.9896 (0.1648)	0.9998 (0.2459)	$\tilde{\nu}_0$	0.1174 (0.1305)	0.0390 (0.0863)	0.2327 (0.0857)
ν_1	2.1313 (0.1346)	3.6763 (0.2959)	4.1018 (0.8327)	$\tilde{\nu}_1$	-0.0403 (0.1339)	0.1693 (0.1897)	-0.8568 (0.4269)
ν_2	-2.0959 (0.1488)	-3.4849 (0.2944)	-3.7677 (0.8411)	$\tilde{\nu}_2$	0.0195 (0.1419)	-0.1225 (0.2011)	0.9965 (0.4285)
$\nu_{c,1}$	0.1749 (0.0119)	0.2815 (0.0304)	0.3137 (0.0892)	$\tilde{\nu}_{c,1}$	-0.0052 (0.0138)	0.0124 (0.0216)	-0.1045 (0.0454)
$\nu_{c,2}$	0.0194 (0.0095)	0.0679 (0.0109)	0.1277 (0.0145)	$\tilde{\nu}_{c,2}$	-0.0075 (0.0098)	0.0192 (0.0132)	0.0492 (0.0150)
$\nu_{s,1}$	0.0307 (0.0043)	0.0475 (0.0102)	0.0440 (0.0273)	$\tilde{\nu}_{s,1}$	0.0006 (0.0052)	0.0132 (0.0076)	-0.0258 (0.0156)
$\nu_{s,2}$	0.0099 (0.0045)	0.0345 (0.0080)	0.0676 (0.0115)	$\tilde{\nu}_{s,2}$	-0.0008 (0.0039)	0.0091 (0.0126)	0.0523 (0.0130)
Total No. of durations	23218	11327	7292				
log-likelihood	-28526	-14145	-9297				

Parameter estimates (standard errors in parentheses) for the log-ACD-GARCH model (with seasonality components s and \tilde{s}) in Eq. (21). The data-generating process (DGP) follows the Heston model in Eq. (25). Durations are obtained with the K -adaptive thresholds for $K = 78, 39$, and 26 . The log-ACD-GARCH- s, \tilde{s} model is estimated via MLE with the assumption that the white noise $u_{d,i}$ follows a standard normal distribution. We select the lags (1, 1) for both the ACD and GARCH components, and employ the second-order flexible-Fourier-form specifications in Eq. (22) with $Q = 2$ for both seasonality terms s and \tilde{s} .

3.3 Spot Volatility Estimation

In this section, we utilize the results in Section 2.1 to estimate the intraday spot volatility with a local IV estimator:

$$\hat{\sigma}_t = \sqrt{\frac{1}{\Delta} \widehat{V}(t - \Delta, t)}, \quad (29)$$

where $\widehat{V}(t - \Delta, t)$ is defined in Eq. (10), and can be estimated with the conditional CDFs of price durations with both the left and right correction. Fig. 3 illustrates an example of spot volatility estimation for each equidistant intervals with $\Delta = 5$ minutes, 30 minutes, and one hour. The solid line represents a simulated path of the spot volatility process σ from the Heston model in Eq. (25). All annualized spot volatility estimates are calculated from Eq. (29) based on the log-

ACD-GARCH- s , \tilde{s} model in Eq. (21), with durations obtained with the daily K -adaptive thresholds for $K = 78$.

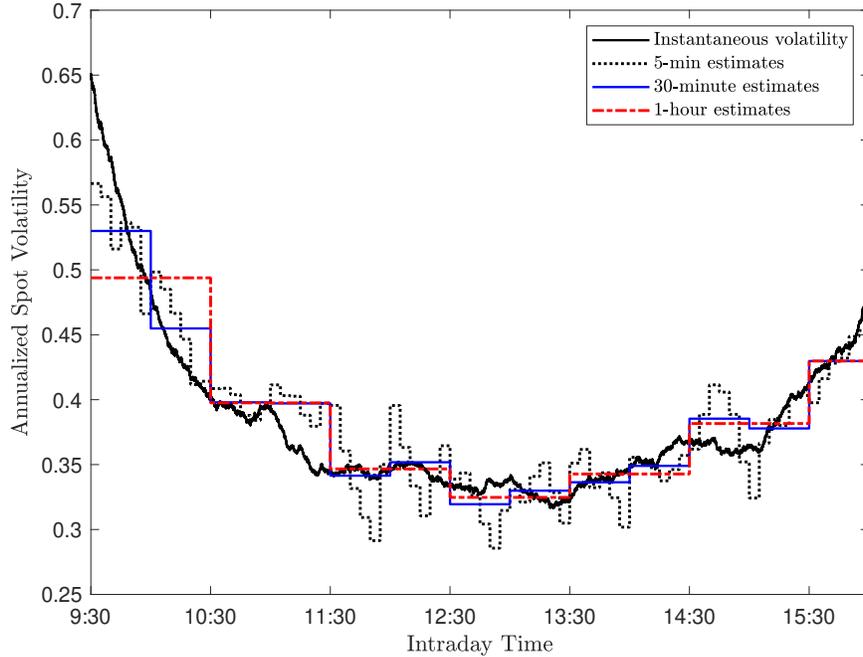


Figure 3: Annualized spot volatility estimates for each equidistant intervals with $\Delta = 5$ minutes, 30 minutes, and one hour. The DGP follows the Heston model in Eq. (25). Durations are obtained with the K -adaptive threshold with $K = 78$. The duration-based spot volatility estimator is constructed based on the log-ACD-GARCH- s , \tilde{s} model in Eq. (21).

Throughout this section, we evaluate the finite-sample performance with the root-mean-square error (RMSE), i.e., for each spot volatility estimate $\hat{\sigma}_i$ over the i -th interval,

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_i - \sigma_i)^2}, \quad (30)$$

where N represents the total number of intervals (number of intervals per day \times 300 days), and the “true value” σ_i is the local average of all annualized spot volatilities (tick-level) within each interval.⁴

Table 2 reports the RMSE results for our duration-based spot volatility estimators across equidistant intervals with $\Delta = 5, 10, 15, 30,$ and 60 minutes, based on durations obtained with three different levels of $K = 78, 39,$ and 26 . These results are derived from four parametric duration models: (i) the multiplicative ACD model (Engle and Russell, 1998; Tse and Yang, 2012), (ii) the log-ACD model, (iii) the log-ACD-GARCH model, and (iv) the log-ACD-GARCH- s , \tilde{s} model in Eq. (21). For model estimation, we follow Section 3.2 to assume that $(u_{d,i})$ follows a standard Gaussian white noise process for all three log-ACD models, while the residuals for the multiplicative ACD model are assumed to be exponentially distributed, i.e., $\text{Exp}(1)$.

The RMSE results for all out-of-sample days are reported in Panel A. We observe that the spot

Table 2: Monte Carlo RMSEs for duration-based spot volatility estimation

Panel A: All Out-of-Sample Days (300 Days)												
Interval	(i) ACD			(ii) log-ACD			(iii) log-ACD-GARCH			(iv) log-ACD-GARCH- s, \bar{s}		
	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26
	5 min	5.55	6.64	7.31	5.39	6.27	6.58	5.39	6.27	6.58	4.57	5.23
10 min	5.08	6.06	6.57	5.15	6.17	6.52	5.15	6.17	6.52	4.32	5.14	5.53
15 min	4.89	5.66	6.35	4.98	5.85	6.59	4.98	5.85	6.59	4.17	4.76	5.24
30 min	4.57	5.18	5.69	4.61	5.36	5.96	4.62	5.36	5.96	3.93	4.31	4.92
60 min	4.15	4.76	5.26	4.16	4.85	5.36	4.16	4.85	5.36	3.77	4.17	4.65

Panel B: Days with Jumps (51 Days)												
Interval	(i) ACD			(ii) log-ACD			(iii) log-ACD-GARCH			(iv) log-ACD-GARCH- s, \bar{s}		
	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26
	5 min	5.73	6.80	7.43	5.48	6.41	6.74	5.48	6.41	6.74	4.52	5.15
10 min	5.11	6.20	6.69	5.09	6.32	6.73	5.09	6.32	6.73	4.05	5.05	5.44
15 min	4.93	5.55	6.09	4.94	5.80	6.40	4.94	5.80	6.40	3.92	4.40	5.01
30 min	4.56	5.18	5.77	4.54	5.45	6.14	4.54	5.45	6.14	3.66	4.13	4.76
60 min	4.08	4.71	5.06	4.04	4.87	5.36	4.05	4.87	5.35	3.48	3.97	4.30

Panel C: Days without Jumps (249 Days)												
Interval	(i) ACD			(ii) log-ACD			(iii) log-ACD-GARCH			(iv) log-ACD-GARCH- s, \bar{s}		
	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26
	5 min	5.52	6.61	7.28	5.37	6.24	6.55	5.38	6.24	6.55	4.58	5.24
10 min	5.08	6.02	6.55	5.16	6.14	6.48	5.17	6.14	6.48	4.37	5.16	5.55
15 min	4.89	5.68	6.41	4.99	5.87	6.63	5.00	5.87	6.63	4.22	4.83	5.34
30 min	4.57	5.18	5.67	4.63	5.34	5.92	4.63	5.34	5.92	3.98	4.34	4.96
60 min	4.17	4.78	5.30	4.18	4.84	5.36	4.19	4.84	5.36	3.83	4.21	4.72

RMSE results ($\times 10^2$) for the duration-based spot volatility estimators based on four parametric duration models: (i) the multiplicative ACD model, (ii) the log-ACD model, (iii) the log-ACD-GARCH model, and (iv) the log-ACD-GARCH- s, \bar{s} model in Eq. (21). The DGP follows the Heston model in Eq. (25). Durations are obtained with the K -adaptive thresholds for $K = 78, 39,$ and 26 . The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each equidistant interval.

volatility estimators based on log-ACD models tend to achieve better finite-sample performance than the ACD-based counterpart, particularly for shorter interval lengths. While the inclusion of GARCH components offers little improvement, the inclusion of seasonality components significantly reduces RMSEs across all interval lengths. This underscores the importance of periodic patterns in intraday durations to improve the accuracy of local volatility estimation, and also highlights the flexibility of parametric structures in incorporating additional information.

Table 2 also compares the RMSE results between days with jumps (Panel B) and days without jumps (Panel C). For relatively small local intervals, the RMSEs on days with jumps are slightly higher than those on days with only continuous price movements. However, the difference is minimal and can even be reversed by finite-sample bias as the interval length increases. As highlighted in Remark 1, the inherent robustness of duration-based methods ensures the reliability of our volatility estimator in the presence of relatively large and sudden price changes. To further assess the sensitivity of our duration-based spot volatility estimator to price jumps, we insert a single jump randomly between 12:55 and 13:00 on the same day illustrated in Fig. 3. As demonstrated in

Fig. 4, such a discontinuous shift in the price process has almost no impact on our duration-based estimator but can significantly bias the localized RV over the block where the jump occurs.

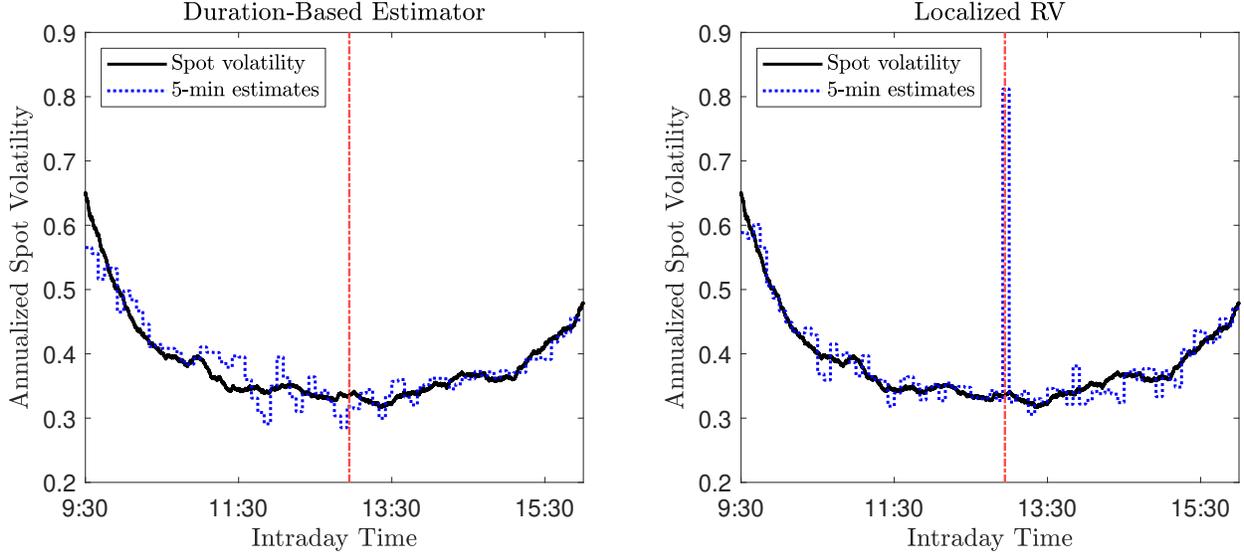


Figure 4: Comparisons of 5-minute spot volatility estimates in the presence of jumps. The DGP follows the Heston model in Eq. (25). Durations are obtained with the K -adaptive threshold for $K = 78$. The duration-based spot volatility estimator is constructed based on the log-ACD-GARCH- s, \tilde{s} model in Eq. (21). The localized RV estimator are constructed from all 5-second returns within each block.

To evaluate the finite-sample performance of our spot volatility estimator, we compare it with some conventional benchmarks. The first one we consider is the classical spot volatility estimator introduced by Foster and Nelson (1996), which is a localized version of RV (Barndorff-Nielsen and Shephard, 2002; Andersen et al., 2003a). Suppose there are k_n price observations with the lag Δ_n over the block $[t - \Delta, t]$, i.e., $\Delta = k_n \Delta_n$. The localized RV estimator is defined as

$$\hat{\sigma}_t = \sqrt{\frac{1}{\Delta} \text{RV}_{\Delta, k_n}}, \quad \text{RV}_{\Delta, k_n} = \sum_{i=1}^{k_n} r_i^2, \quad (31)$$

where $r_i = X_{t-\Delta+i\Delta_n} - X_{t-\Delta+(i-1)\Delta_n}$ is the i -th return between consecutive observations.

In addition to the localized RV, we consider localized versions of two jump-robust realized measures, i.e., the realized bipower variation (BV) of Barndorff-Nielsen and Shephard (2004) and the truncated realized volatility (TRV) of Mancini (2009):⁵

$$\text{BV}_{\Delta, k_n} = \frac{\pi}{2} \frac{k_n}{k_n - 1} \sum_{i=2}^{k_n} |r_i| |r_{i-1}|, \quad \text{TRV}_{\Delta, k_n} = \sum_{i=1}^{k_n} r_i^2 \mathbb{1}_{\{|r_i| \leq \zeta \Delta_n^{\varpi}\}}, \quad (32)$$

with the truncation parameters $\varpi = 0.49$ and ζ determined with the data-adaptive method of Andersen et al. (2023). Specifically, ζ is defined as

$$\zeta = C_\zeta \sqrt{\text{MedRV}_d}, \quad (33)$$

where we select $C_\zeta = 3$, and MedRV is the daily median RV estimator of Andersen et al. (2012), constructed from all one-minute returns within day d .

Panel A of Table 3 presents the RMSE results for various return-based spot volatility estimators for all out-of-sample days. All these estimators are constructed from more granular returns (1-second, 5-second, 30-second, and 1-minute) within each block of Δ . In the absence of market microstructure noise, the second-by-second “tick-level” observations provides the most precise estimates of local volatility for each “not-too-finely” sampled interval. Among the benchmarks, the localized RV estimator demonstrates the largest RMSE values overall. In contrast, the localized BV and TRV achieve significantly lower RMSEs. The breakdown in Panel B and C demonstrates the impact of jumps on these estimators. On days without jumps (Panel C), the localized RV obtains the smallest RMSEs, which is consistent with well-established theoretical results in the literature. However, on days with jumps (Panel B), the localized BV and TRV outperform RV by avoiding drastically inflated RMSEs.

Table 3: Monte Carlo RMSEs for return-based spot volatility estimation

Panel A: All Out-of-Sample Days (300 Days)												
Interval	localized RV				localized BV				localized TRV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	11.05	11.49	13.87	16.23	1.93	4.16	10.21	14.38	2.10	3.71	8.35	11.73
10 min	10.53	10.76	12.14	13.57	1.39	2.95	7.42	10.31	1.81	2.84	6.06	8.46
15 min	10.19	10.36	11.34	12.37	1.16	2.42	6.14	8.44	1.69	2.48	5.03	6.99
30 min	9.53	9.62	10.16	10.74	0.85	1.74	4.43	6.02	1.51	2.03	3.75	5.10
60 min	8.39	8.46	8.83	9.19	0.63	1.34	3.43	4.59	1.30	1.68	2.98	3.99
Panel B: Days with Jumps (51 Days)												
Interval	localized RV				localized BV				localized TRV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	26.57	26.78	27.96	29.27	2.60	4.97	12.31	15.91	2.04	3.75	8.48	12.11
10 min	25.43	25.54	26.18	26.94	1.95	3.69	9.21	11.85	1.72	2.90	6.08	8.87
15 min	24.64	24.73	25.19	25.71	1.66	3.07	7.92	9.89	1.60	2.53	5.01	7.30
30 min	23.07	23.12	23.34	23.62	1.24	2.25	5.89	7.28	1.42	2.03	3.68	5.24
60 min	20.33	20.37	20.53	20.63	0.90	1.71	4.70	5.47	1.21	1.66	2.99	3.96
Panel C: Days without Jumps (249 Days)												
Interval	localized RV				localized BV				localized TRV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	1.55	3.48	8.47	11.91	1.77	3.98	9.80	14.04	2.12	3.71	8.32	11.65
10 min	1.10	2.44	6.11	8.56	1.25	2.78	6.99	9.96	1.82	2.83	6.05	8.38
15 min	0.90	1.99	5.00	6.99	1.03	2.26	5.70	8.11	1.71	2.47	5.04	6.92
30 min	0.65	1.43	3.58	4.99	0.74	1.62	4.07	5.73	1.53	2.03	3.76	5.07
60 min	0.50	1.10	2.75	3.82	0.56	1.25	3.11	4.39	1.31	1.69	2.88	4.00

RMSE results ($\times 10^2$) for the localized return-based estimators based on 1-second, 5-second, 30-second, and 1-minute returns. The DGP follows the Heston model in Eq. (25). The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each equidistant interval.

Moreover, we observe from Table 3 that the accuracy of return-based volatility estimation depends significantly on the granularity of local observations. The localized BV and TRV constructed from 1-second and 5-second returns outperform our duration-based estimator in Table 2 in terms of RMSE. However, in real financial markets, very fine sampling introduces the challenge of market

microstructure noise. The noise captures various idiosyncrasies inherent in the trading process, such as bid-ask bounces and the discreteness of price changes, which becomes increasingly dominant and distorts the volatility estimates as the sampling frequency increases.

Market Microstructure Noise. To examine the impact of market microstructure noise on the finite-sample performance of both the duration-based and return-based estimators, we incorporate a noise component into our simulation framework. Specifically, we augment the Heston model in Eq. (25) with an additive heterogeneous Gaussian noise term for the simulation of all second-by-second transactions, which follows the simulation specifications in Aït-Sahalia et al. (2012) and Christensen et al. (2022):

$$Y_i = X_i + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \omega_i^2), \quad \text{where } \omega_i = \gamma \sqrt{\frac{\sigma_i^2}{n}}, \quad (34)$$

for $i = 0, 1, \dots, n$, where the noise-to-volatility ratio is set to $\gamma = 1.5$.⁶

Table 4 reports the RMSE results for our duration-based spot volatility estimators in the presence of market microstructure noise. While the RMSEs show slight inflation compared to the results in Table 2, the performance remains largely comparable to the no-noise scenario. Consistent with the noise-free results, the log-ACD-GARCH- s, \tilde{s} model continues to outperform the other models in terms of RMSE. As discussed in Section 2, the adaptive choice of sampling thresholds provides a natural control for the daily expected sampling frequencies, which parallels the common use of sparsely sampled data in both financial econometrics literature and practice (typically equidistant in calendar time) to mitigate the impact of market microstructure noise (Aït-Sahalia et al., 2005; Liu et al., 2015). The results in Table 4 confirm the effective protection against market imperfections provided by the adaptive PDS algorithm, and demonstrate the robustness of duration-based volatility estimators under noise contamination.

The RMSE results in Table 5 illustrate the impact of market microstructure noise on localized return-based estimators. Across all panels, the presence of noise significantly inflates the RMSEs for estimators constructed from local observations at finer intervals, such as 1-second and 5-second returns. As the sampling interval increases, e.g., from 1-second to 1-minute returns, the influence of noise diminishes, and thus the RMSEs decrease across all estimators. However, this reduction in noise sensitivity comes at the cost of losing granular information, which leaves the performance uniformly inferior to that of the duration-based estimators reported in Table 4.

Furthermore, we consider localized versions of noise-corrected return-based estimators that leverage all noise-contaminated tick-level transactions, such as the pre-averaged realized volatility (PRV) and bipower variation (PBV), as detailed in Jacod et al. (2009), Podolskij and Vetter (2009), and Christensen et al. (2025). These estimators are designed to retain granular information but also mitigate the impact of market microstructure noise. In our simulations, the pre-averaging window for both localized PRV and PBV is set to $\lceil \theta \sqrt{k_n} \rceil$, where $\theta \in \{0.3, 0.5, 0.7, 1.0\}$, and k_n counts the total number of “tick-level” second-by-second returns within each interval of Δ .

Table 4: Monte Carlo RMSEs for duration-based spot volatility estimation under market microstructure noise

Panel A: All Out-of-Sample Days (300 Days)												
Interval	(i) ACD			(ii) log-ACD			(iii) log-ACD-GARCH			(iv) log-ACD-GARCH- s, \bar{s}		
	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26
	5 min	9.83	10.00	10.10	9.57	9.61	9.70	9.56	9.61	9.69	8.52	8.28
10 min	9.49	9.54	9.60	9.31	9.49	9.74	9.30	9.49	9.73	8.29	8.18	8.20
15 min	9.38	9.23	9.34	9.19	9.20	9.63	9.18	9.62	9.62	8.19	7.85	8.10
30 min	9.19	8.97	9.00	8.93	8.90	9.30	8.93	8.90	9.29	8.03	7.63	7.81
60 min	8.96	8.69	8.56	8.65	8.50	8.69	8.64	8.50	8.69	8.04	7.53	7.44
Panel B: Days with Jumps (51 Days)												
Interval	(i) ACD			(ii) log-ACD			(iii) log-ACD-GARCH			(iv) log-ACD-GARCH- s, \bar{s}		
	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26
	5 min	8.83	9.26	9.53	8.65	8.92	9.05	8.62	8.89	9.05	7.53	7.50
10 min	8.85	8.77	8.92	8.44	8.94	9.09	8.42	8.91	9.09	7.35	7.55	7.38
15 min	8.41	8.60	8.77	8.35	8.71	8.99	8.33	8.68	8.98	7.29	7.31	7.30
30 min	8.25	8.30	8.62	8.10	8.35	8.96	8.09	8.33	8.96	7.16	7.04	7.39
60 min	7.85	7.83	7.83	7.63	7.74	8.02	7.61	7.73	8.02	6.95	6.70	6.60
Panel C: Days without Jumps (249 Days)												
Interval	(i) ACD			(ii) log-ACD			(iii) log-ACD-GARCH			(iv) log-ACD-GARCH- s, \bar{s}		
	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26	$K = 78$	39	26
	5 min	10.02	10.15	10.21	9.74	9.74	9.83	9.74	9.74	9.82	8.71	8.43
10 min	9.68	9.69	9.74	9.48	9.60	9.86	9.48	9.60	9.86	8.46	8.30	8.36
15 min	9.56	9.36	9.45	9.35	9.30	9.75	9.35	9.30	9.74	8.36	7.96	8.26
30 min	9.37	9.10	9.08	9.09	9.01	9.37	9.09	9.01	9.36	8.20	7.75	7.89
60 min	9.17	8.85	8.70	8.84	8.65	8.82	8.84	8.65	8.82	8.25	7.69	7.60

RMSE results ($\times 10^2$) for the duration-based spot volatility estimators based on four parametric duration models: (i) the multiplicative ACD model, (ii) the log-ACD model, (iii) the log-ACD-GARCH model, and (iv) the log-ACD-GARCH- s, \bar{s} model in Eq. (21). The DGP follows the Heston model in Eq. (25) with an additive heterogeneous Gaussian noise term in Eq. (34). Durations are obtained with the K -adaptive thresholds for $K = 78, 39,$ and 26 . The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each equidistant interval.

Table 6 presents the RMSE results for both localized pre-averaged estimators. We observe that PRV and PBV exhibit stable finite-sample performance across different pre-averaging window sizes. Compared to the results based on second-by-second local observations in Table 4, the pre-averaging approach proposed by Jacod et al. (2009) significantly mitigates the impact of market microstructure noise. Moreover, consistent with the results for localized RV and BV in Table 3, the RMSEs for PRV are drastically inflated on days with jumps (Panel B). In contrast, PBV demonstrates excellent robustness to price jumps, with consistently lower RMSEs across all interval lengths and θ choices. More importantly, the comparison of RMSEs between Tables 4 and 6 shows that our duration-based estimator performs comparably, with slightly better results than the more granular PRV and PBV. This underscores the reliability and practical effectiveness of our new duration-based method.

Table 5: Monte Carlo RMSEs for return-based spot volatility estimation under market microstructure noise

Panel A: All Out-of-Sample Days (300 Days)												
Interval	localized RV				localized BV				localized TRV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	55.95	21.66	16.69	18.34	58.67	19.55	13.97	16.88	26.39	15.31	12.08	14.32
10 min	55.86	21.24	15.16	15.92	58.62	19.21	11.81	13.39	26.39	15.08	10.51	11.68
15 min	55.81	21.03	14.50	14.89	58.60	19.09	10.98	11.90	26.34	15.00	9.90	10.60
30 min	55.73	20.69	13.56	13.56	58.58	18.96	10.08	10.31	26.28	14.89	9.27	9.43
60 min	55.85	20.32	12.68	12.42	58.84	19.01	9.74	9.54	26.00	14.83	9.04	8.95
Panel B: Days with Jumps (51 Days)												
Interval	localized RV				localized BV				localized TRV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	59.63	31.96	29.20	30.02	57.09	18.81	14.60	17.59	25.71	14.08	11.65	14.16
10 min	59.28	30.99	27.47	27.74	57.03	18.38	12.25	14.15	25.17	13.83	9.99	11.37
15 min	59.08	30.40	26.51	26.60	57.01	18.24	11.17	12.31	25.15	13.74	9.24	10.22
30 min	58.73	29.31	24.83	24.68	56.99	18.04	10.07	10.67	25.09	13.61	8.58	9.03
60 min	58.28	27.49	22.25	21.91	57.25	17.98	9.50	9.46	24.78	13.53	8.30	8.39
Panel C: Days without Jumps (249 Days)												
Interval	localized RV				localized BV				localized TRV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	55.17	18.87	12.69	14.86	58.98	19.70	13.84	16.73	26.63	15.56	12.17	14.35
10 min	55.13	18.62	11.06	12.16	58.94	19.38	11.72	13.23	26.59	15.33	10.62	11.74
15 min	55.11	18.53	10.45	11.06	58.92	19.27	10.94	11.82	26.57	15.25	10.03	10.67
30 min	55.10	18.44	9.79	9.84	58.90	19.15	10.08	10.24	26.51	15.13	9.40	9.51
60 min	55.34	18.51	9.60	9.36	59.16	19.21	9.79	9.56	26.24	15.08	9.19	9.06

RMSE results ($\times 10^2$) for the localized return-based estimators based on 1-second, 5-second, 30-second, and 1-minute returns. The DGP follows the Heston model in Eq. (25) with an additive heterogeneous Gaussian noise term in Eq. (34). The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each equidistant interval.

3.4 Integrated Variance Estimation

Similar to the localized IV estimator used in Section 3.3, the theoretical results in Section 2.1 can also be employed to estimate the daily IV:

$$\widehat{V}_d = \widehat{V}(d-1, d) = \delta_d^2 \left[\sum_{i=1}^{N_d} \widehat{G}_{d,i}(x_{d,i}) + \widehat{G}_{d,N_d+1}(d - \tau_{d,N_d}) \right], \quad (35)$$

where the second term in the summand represents the end-of-day correction between the last price event and the market closing time. To estimate the daily IV, we compute all durations between price events over a 300-day period with the predicted K -adaptive thresholds for $K = 78, 39$, and 26 . The conditional CDFs of durations are then used to obtain the IV estimates.

For the comparative analysis, we consider the daily RV, BV, TRV, MinRV, and MedRV as benchmarks. All selected return-based IV estimators are constructed from 5, 10, and 15-minute sampled returns, corresponding to the sampling frequencies of the adaptive PDS with the chosen values of K . For the truncation threshold of TRV, we follow the instruction in Section 3.3.

Table 7 presents the RMSE results for both duration-based and return-based IV estimators in the absence of market microstructure noise. Panel A compares their overall performance across

Table 6: Monte Carlo RMSEs for localized PRV and PBV

Panel A: All Out-of-Sample Days (300 Days)								
Interval	localized PRV				localized PBV			
	$\theta = 0.3$	0.5	0.7	1.0	$\theta = 0.3$	0.5	0.7	1.0
5 min	14.94	14.09	14.38	14.57	11.51	9.42	9.34	9.68
10 min	13.14	13.59	13.64	13.94	8.53	8.65	8.89	9.12
15 min	12.83	13.13	13.28	13.53	8.24	8.42	8.61	8.83
30 min	12.35	12.54	12.70	12.81	8.04	8.22	8.43	8.63
60 min	12.36	12.39	12.46	12.46	9.11	9.16	9.28	9.43
Panel B: Days with Jumps (51 Days)								
Interval	localized PRV				localized PBV			
	$\theta = 0.3$	0.5	0.7	1.0	$\theta = 0.3$	0.5	0.7	1.0
5 min	27.28	27.27	28.34	28.52	12.16	9.85	9.66	10.13
10 min	25.92	26.89	26.91	27.30	8.68	8.77	9.06	9.48
15 min	25.35	25.82	26.08	26.31	8.30	8.49	8.74	9.09
30 min	24.08	24.35	24.51	24.51	7.85	8.15	8.36	8.63
60 min	22.16	22.13	22.07	22.03	9.00	9.08	9.24	9.31
Panel C: Days without Jumps (249 Days)								
Interval	localized PRV				localized PBV			
	$\theta = 0.3$	0.5	0.7	1.0	$\theta = 0.3$	0.5	0.7	1.0
5 min	10.80	9.05	9.02	9.45	11.37	9.33	9.27	9.59
10 min	8.38	8.62	8.79	9.19	8.51	8.64	8.75	9.15
15 min	8.18	8.37	8.57	8.83	8.08	8.21	8.31	8.68
30 min	8.05	8.25	8.43	8.63	8.08	8.23	8.42	8.61
60 min	9.11	9.15	9.26	9.35	8.91	9.15	9.19	9.35

RMSE results ($\times 10^2$) for the localized pre-averaged realized volatility (PRV) and bipower variation (PBV) constructed from all second-by-second observations. The DGP follows the Heston model in Eq. (25) with an additive heterogeneous Gaussian noise term in Eq. (34). The pre-averaging window is set to $\lceil \theta \sqrt{k_n} \rceil$. The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each equidistant interval.

all out-of-sample days, which demonstrates that the duration-based estimators, derived from four different duration models, consistently outperform return-based estimators with lower RMSEs across all selected sampling frequencies. On days with jumps (Panel B), the duration-based estimators exhibit superior robustness. While jump-robust return-based estimators, such as BV and TRV, show some improvement over the standard RV, their RMSEs remain higher than those of the duration-based estimators.

Furthermore, we account for the impact of market microstructure noise by incorporating the additive heterogeneous Gaussian noise term in Eq. (34). In Table 8, we include all duration-based and return-based IV estimators from Table 7, while also adding pre-averaged estimators, specifically PRV and PBV. Both of them are constructed from the “tick-level” second-by-second and “finely sampled” 5-second price observations. Compared with the noise-free results in Table 7, the presence of noise only slightly deteriorates the finite-sample performance of both duration-based and return-based estimators based on sparse sampling. Notably, our duration-based estimators demonstrate superior robustness and achieve more stable accuracy across various choices of duration models and sampling frequencies.

While the pre-averaged estimators are specifically designed to mitigate the impact of noise and to utilize the most granular information, our duration-based estimators achieve RMSE results

Table 7: Monte Carlo RMSEs for IV estimation

Panel A: All Out-of-Sample Days (300 Days)									
K	Duration-based estimators				Return-based estimators				
	(i)	(ii)	(iii)	(iv)	RV	BV	TRV	MinRV	MedRV
78	2.60	2.06	2.06	2.05	6.37	3.01	2.78	3.42	2.90
39	2.19	2.21	2.21	2.13	7.07	4.15	3.97	5.09	4.23
26	2.41	2.40	2.29	2.25	7.57	4.84	4.38	5.31	4.90
Panel B: Days with Jumps (51 Days)									
K	Duration-based estimators				Return-based estimators				
	(i)	(ii)	(iii)	(iv)	RV	BV	TRV	MinRV	MedRV
78	2.59	1.90	1.90	1.88	14.51	3.94	3.99	3.38	2.97
39	2.03	2.13	2.13	1.99	15.37	4.91	4.41	4.62	3.86
26	2.26	2.37	2.30	2.08	16.31	7.29	5.77	5.97	6.29
Panel C: Days without Jumps (249 Days)									
K	Duration-based estimators				Return-based estimators				
	(i)	(ii)	(iii)	(iv)	RV	BV	TRV	MinRV	MedRV
78	2.61	2.09	2.09	2.60	2.41	2.78	2.81	3.42	2.88
39	2.22	2.23	2.23	2.16	3.43	3.98	3.87	5.18	4.31
26	2.45	2.40	2.29	2.26	3.80	4.16	4.04	5.17	4.56

RMSE results for both duration-based and return-based IV estimators in the absence of market microstructure noise. The duration-based estimators are derived from four parametric duration models: (i) the multiplicative ACD model, (ii) the log-ACD model, (iii) the log-ACD-GARCH model, and (iv) the log-ACD-GARCH- s, \tilde{s} model in Eq. (21). The DGP follows the Heston model in Eq. (25). Durations are obtained with the K -adaptive thresholds for $K = 78, 39,$ and 26 . For comparison, all selected return-based IV estimators are constructed from 5, 10, and 15-minute sampled returns, which maintain the same sampling frequencies as the adaptive PDS. The RMSE results are calculated based on based on the annualized IV estimates for 300 out-of-sample days.

comparable to those of PBV and outperform PRV, particularly on days with jumps. These results highlight the robustness and effectiveness of our duration-based method as a reliable alternative to return-based approaches for high-frequency IV estimation.

4 Empirical Analysis

In this section, we first utilize our duration-based IV estimator as the basis for daily volatility forecasting under the HAR framework for the SPDR S&P 500 ETF Trust (SPY), which is the best-recognized and oldest U.S. listed ETF and by far the most widely traded S&P 500 ETF. Subsequently, we employ our duration-based spot volatility estimator to assess how intraday volatility responds to some specific macroeconomic events, such as FOMC news announcements.

4.1 Overview

We obtain all high-frequency transaction records of SPY from the daily Trade and Quote (TAQ) dataset, with the sample period ranging from January 2, 2014 to December 30, 2022. The tick-by-tick transactions are timestamped in milliseconds until mid-2015 and in microseconds since then.⁷ As is standard in empirical research with TAQ data, we use the filters as in [Barndorff-Nielsen et al. \(2009\)](#) to eliminate data errors, remove all transactions in the original record that are later corrected, cancelled or otherwise invalidated. In addition, we remove all trading days with an early market

Table 8: Monte Carlo RMSEs for IV estimation under market microstructure noise

Panel A: All Out-of-Sample Days (300 Days)												
K	Duration-based estimators				Return-based estimators					K	PRV	PBV
	(i)	(ii)	(iii)	(iv)	RV	BV	TRV	MinRV	MedRV			
78	2.85	2.65	2.64	2.56	6.71	3.49	3.04	3.91	3.41	23,400	6.43	2.24
39	2.86	2.69	2.68	2.66	7.21	4.69	3.99	5.32	4.71	4,680	6.64	2.83
26	3.01	2.82	2.81	2.76	8.13	5.43	4.79	6.34	5.60	–	–	–
Panel B: Days with Jumps (51 Days)												
K	Duration-based estimators				Return-based estimators					K	PRV	PBV
	(i)	(ii)	(iii)	(iv)	RV	BV	TRV	MinRV	MedRV			
78	2.86	2.73	2.70	2.64	15.05	4.58	3.07	3.59	3.17	23,400	15.00	2.66
39	2.86	2.70	2.68	2.65	15.45	6.99	4.49	6.71	6.04	4,680	15.14	3.57
26	2.99	2.81	2.80	2.72	17.04	7.49	5.34	7.08	6.44	–	–	–
Panel C: Days without Jumps (249 Days)												
K	Duration-based estimators				Return-based estimators					K	PRV	PBV
	(i)	(ii)	(iii)	(iv)	RV	BV	TRV	MinRV	MedRV			
78	2.84	2.64	2.62	2.54	2.81	3.23	3.03	3.97	3.46	23,400	1.96	2.14
39	2.86	2.69	2.68	2.66	3.71	4.06	3.89	4.98	4.39	4,680	2.50	2.65
26	3.02	2.83	2.82	2.76	4.49	4.90	4.67	6.18	5.42	–	–	–

RMSE results for both duration-based and return-based IV estimators in the presence of market microstructure noise. The duration-based estimators are derived from four parametric duration models: (i) the multiplicative ACD model, (ii) the log-ACD model, (iii) the log-ACD-GARCH model, and (iv) the log-ACD-GARCH- s, \tilde{s} model in Eq. (21). The DGP follows the Heston model in Eq. (25) with an additive heterogeneous Gaussian noise term in Eq. (34). Durations are obtained with the K -adaptive thresholds for $K = 78, 39,$ and 26 . For comparison, all selected return-based IV estimators are constructed from 5, 10, and 15-minute sampled returns, which maintain the same sampling frequencies as the adaptive PDS. For PRV and PBV, the pre-averaging window is set to $\lceil \theta \sqrt{n} \rceil$ with $\theta = 0.5$. The RMSE results are calculated based on based on the annualized IV estimates for 300 out-of-sample days.

closure, and restrict our sample to transactions between 9:30:00 – 16:00:00 Eastern Time (ET) for all individual stocks.

To obtain the durations with the adaptive PDS, we utilize all trading days from January 2014 to December 2016 as the burn-in period, which allows us to predict the first K -adaptive daily threshold in 2017. Fig. 5 compares the autocorrelations of log-durations (from 2017 to 2022, 1499 days in total) obtained with the predicted daily thresholds (with $K = 78$) to those obtained with a fixed threshold for all days. It is notable that the utilization of daily adaptive thresholds effectively alleviates the long memory observed in log-durations over extended multi-day periods.

As an example, we estimate the log-ACD-GARCH- s, \tilde{s} model in Eq. (21) via MLE with all log-durations from 2017 to 2022. In line with Section 3, we select the lags (1, 1) for both the ACD and GARCH parts, with second-order flexible-Fourier-form specifications for both seasonality terms s and \tilde{s} . The parameter estimates (with standard errors) are reported in Table 9.⁸ All parameters in both ACD and GARCH components are statistically significant. Both the conditional means and variances demonstrate strong positive autocorrelation. For the intraday seasonality terms s and \tilde{s} , most of the estimated parameters in the first sine-cosine summand ($\nu_{c,1}, \nu_{s,1}, \tilde{\nu}_{c,1}$) are significant. Fig. 6 illustrates the intraday seasonality for the conditional mean and variance of durations (with $K = 78$), respectively.

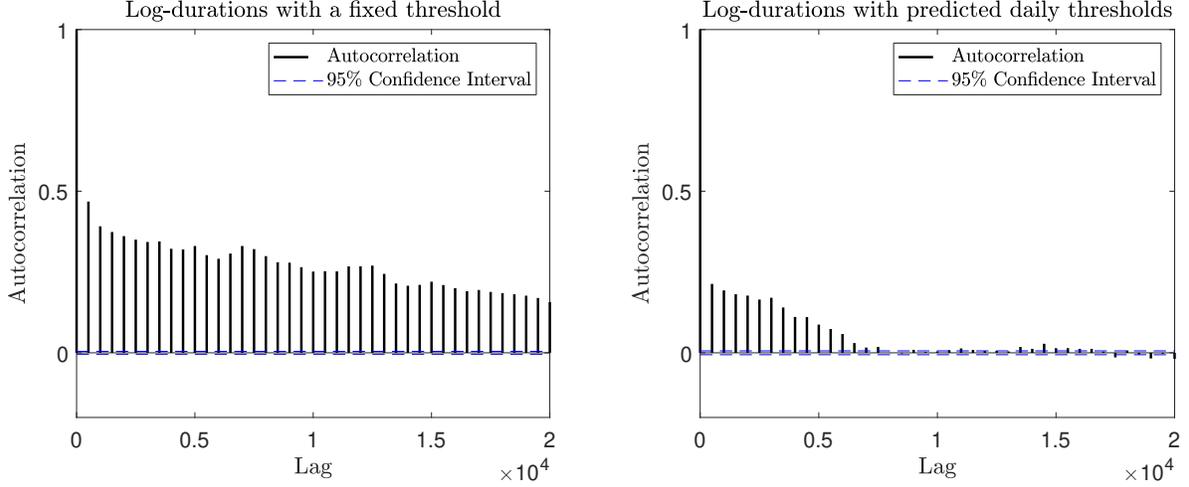


Figure 5: Correlograms of log-durations obtained with (a) a fixed threshold and (b) the predicted daily thresholds with $K = 78$. We choose the fixed threshold as the mean value of ex-post K -adaptive thresholds in Eq. (18) for all trading days.

Table 9: Parameter estimates of the log-ACD-GARCH- s, \bar{s} model

Parameters	$K = 78$	$K = 39$	$K = 26$	Parameters	$K = 78$	$K = 39$	$K = 26$
ϕ_1	0.9774 (0.0007)	0.9185 (0.0142)	0.9271 (0.0032)	α_1	0.0924 (0.0029)	0.0945 (0.0091)	0.0514 (0.0037)
θ_1	-0.7668 (0.0028)	-0.6674 (0.0142)	-0.6558 (0.0073)	β_1	0.6883 (0.0127)	0.5779 (0.1317)	0.8511 (0.0146)
ν_0	-0.1063 (0.0126)	0.0266 (0.1591)	-0.1419 (0.0214)	$\tilde{\nu}_0$	0.4172 (0.0273)	0.5682 (0.1504)	0.2273 (0.0164)
ν_1	1.4692 (0.0716)	2.7307 (0.6385)	3.5099 (0.0890)	$\tilde{\nu}_1$	-0.9631 (0.1150)	-0.9972 (0.1034)	-0.8827 (0.0493)
ν_2	-1.5433 (0.0721)	-2.7181 (0.7036)	-3.4330 (0.0971)	$\tilde{\nu}_2$	0.9999 (0.1130)	0.9994 (0.1121)	0.9996 (0.0514)
$\nu_{c,1}$	0.1306 (0.0076)	0.2147 (0.0754)	0.2806 (0.0098)	$\tilde{\nu}_{c,1}$	-0.1166 (0.0111)	-0.1599 (0.0431)	-0.1104 (0.0062)
$\nu_{c,2}$	0.0020 (0.0020)	0.0501 (0.0336)	0.0886 (0.0066)	$\tilde{\nu}_{c,2}$	0.0022 (0.0042)	-0.0257 (0.0201)	0.0380 (0.0067)
$\nu_{s,1}$	0.0322 (0.0024)	0.0643 (0.0244)	0.0813 (0.0048)	$\tilde{\nu}_{s,1}$	0.0018 (0.0043)	0.0040 (0.0189)	-0.0064 (0.0042)
$\nu_{s,2}$	-0.0029 (0.0017)	0.0204 (0.0153)	0.0318 (0.0051)	$\tilde{\nu}_{s,2}$	0.0273 (0.0034)	0.0620 (0.0329)	0.0339 (0.0063)
Total No. of durations	111716	52015	34079				
log-likelihood	-169462	-79913	-53264				

Parameter estimates (standard errors in parentheses) for the log-ACD-GARCH- s, \bar{s} model in Eq. (21). The white noise $u_{d,i}$ is assumed to follow a standard normal distribution. We select the lags (1, 1) for both the ACD and GARCH components, and utilize the second-order flexible-Fourier-form specifications in Eq. (22) with $Q = 2$ for both seasonality terms s and \bar{s} .

4.2 Daily Volatility Forecasting

We denote the benchmark IV estimator at day d by \widehat{V}_d^* , and \widehat{V}_d represents a generic IV estimator. We define the following moving averages of \widehat{V}_d as:

$$\widehat{V}_d^{(w)} = \frac{1}{5} \sum_{i=1}^5 \widehat{V}_{d-i+1} \quad \text{and} \quad \widehat{V}_d^{(m)} = \frac{1}{22} \sum_{i=1}^{22} \widehat{V}_{d-i+1}, \quad (36)$$

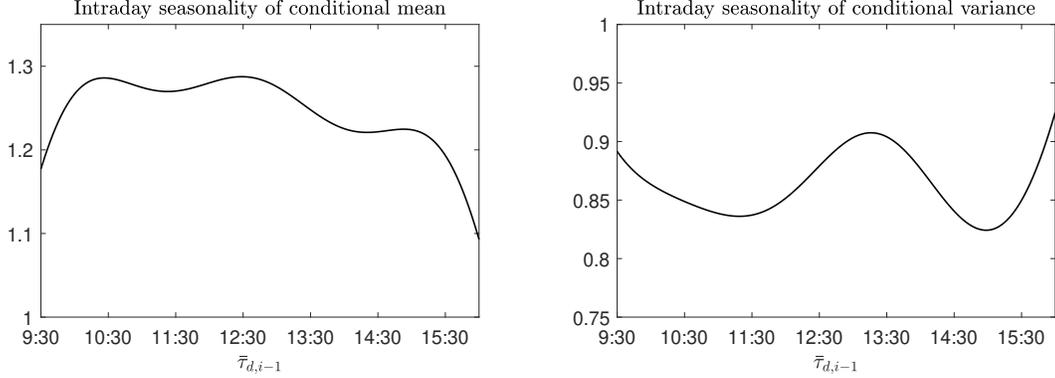


Figure 6: Intraday seasonality $\exp(s_{d,i})$ and $\exp(\tilde{s}_{d,i})$ for the conditional mean and variance of durations (with $K = 78$). The parameter estimates are reported in Table 9.

where $\widehat{V}_d^{(w)}$ represents the one-week average and $\widehat{V}_d^{(m)}$ denotes the one-month average of daily IV estimates, respectively. The standard one-day-ahead HAR model has the following structure:

$$\widehat{V}_d = \omega + \beta^{(d)}\widehat{V}_{d-1} + \beta^{(w)}\widehat{V}_{d-1}^{(w)} + \beta^{(m)}\widehat{V}_{d-1}^{(m)} + \varepsilon_d. \quad (37)$$

As demonstrated in prior empirical studies, incorporating a volatility measure of the continuous component on the right-hand side (RHS) can effectively improve the predictive power of the HAR model for the left-hand side (LHS) target variable \widehat{V}_d^* (Andersen et al., 2007a). In this section, we evaluate the predictive accuracy of HAR models augmented with various IV measures, including our duration-based IV estimators derived from different parametric duration models.

To construct our duration-based IV estimators with the adaptive PDS, we utilize all trading days from January 2014 to December 2016 as the burn-in period, which allows us to forecast the first K -adaptive daily threshold in 2017. The initial in-sample period includes 1000 days from January 3, 2017. Consistent with the Monte Carlo simulations in Section 3, we estimate four parametric duration models, i.e., (i) the multiplicative ACD model, (ii) the log-ACD model, (iii) the log-ACD-GARCH model, and (iv) the log-ACD-GARCH- s, \tilde{s} model in Eq. (21), to obtain all IV estimates in the initial in-sample period, and then forecast the first out-of-sample IV proxy (on the 1001-th day) with the HAR model in Eq. (37). This procedure of in-sample estimation and out-of-sample forecasting is repeated in both rolling-window (RW) and expanding-window (EW) fashion. For the actual values of out-of-sample duration-based IV estimates on subsequent days, we repeatedly re-estimate the duration models with all durations in the corresponding “in-sample plus one day” period.

In addition to the standard HAR model augmented with duration-based estimators and several return-based estimators discussed in Section 3.4, we also consider two important extensions of the original HAR-RV model of Corsi (2009). The first is the quarticity expanded HAR (HARQ) model of Bollerslev et al. (2016). Motivated by the fact that the persistence of RV is influenced by temporal variations in its measurement errors, the HARQ-RV model incorporates a time-varying coefficient

for the previous day’s RV on the RHS. The coefficient is determined by realized quarticity (RQ), which captures the heteroskedasticity in the measurement error:⁹

$$\widehat{V}_d = \omega + \left(\beta^{(d)} + \beta^{(q)} \sqrt{\text{RQ}_{d-1}} \right) \text{RV}_{d-1} + \beta^{(w)} \text{RV}_{d-1}^{(w)} + \beta^{(m)} \text{RV}_{d-1}^{(m)} + \varepsilon_d, \quad (38)$$

where $\text{RQ}_d = 3^{-1} n \sum_{i=1}^n r_{d,i}^4$. Inspired by the intuition that “good” and “bad” volatilities have different effects, the semivariance HAR (SHAR) model of [Patton and Sheppard \(2015\)](#) stands out as another important HAR-RV extension:

$$\widehat{V}_d = \omega + \beta_-^{(d)} \text{RS}_{d-1}^- + \beta_+^{(d)} \text{RS}_{d-1}^+ + \beta^{(w)} \text{RV}_{d-1}^{(w)} + \beta^{(m)} \text{RV}_{d-1}^{(m)} + \varepsilon_d, \quad (39)$$

where the realized semivariances (RS) are introduced by [Barndorff-Nielsen et al. \(2010\)](#):

$$\text{RS}_d^- = \sum_{i=1}^n r_{d,i}^2 \mathbb{1}_{\{r_{d,i} < 0\}} \quad \text{and} \quad \text{RS}_d^+ = \sum_{i=1}^n r_{d,i}^2 \mathbb{1}_{\{r_{d,i} > 0\}}. \quad (40)$$

In this section, we adopt $K = 78$ for the adaptive PDS algorithm, which aligns with the sampling frequency of the commonly used 5-minute CTS. With a uniform threshold applied throughout the day, an important property of PDS is that more samples are generated during periods of higher volatility. An alternative approach for achieving a similar importance-sampling effect adapted to the volatility level is tick-time sampling (TTS), also known as transaction-time sampling. Under TTS, each interval contains a fixed number of trades (ticks) rather than a fixed duration, which makes it convenient to control the sample size ([Andersen et al., 2012](#); [Hautsch and Podolskij, 2013](#)). For our empirical application, we construct the return-based estimators from returns sampled under both CTS and TTS, with a sampling frequency of $K = 78$. Moreover, we construct both PRV and PBV from all available tick-level transaction data with the pre-averaging approach of [Jacod et al. \(2009\)](#).

We evaluate the out-of-sample forecasting performance via two widely used loss functions, i.e., the mean squared error (MSE) and the quasi-likelihood (QLIKE) function:

$$\text{MSE} = \frac{1}{M} \sum_{d=1}^M (\widehat{V}_d^* - \widetilde{V}_d^*)^2 \quad \text{and} \quad \text{QLIKE} = \frac{1}{M} \sum_{d=1}^M \left(\frac{\widehat{V}_d^*}{\widetilde{V}_d^*} - \ln \left(\frac{\widehat{V}_d^*}{\widetilde{V}_d^*} \right) - 1 \right), \quad (41)$$

where \widehat{V}_d^* and \widetilde{V}_d^* denote the ex-post estimate and the forecast of the target IV proxy on day d , respectively, and M represents the total number of out-of-sample days.

Table 10 presents the MSE and QLIKE results for one-day-ahead out-of-sample forecasts of two target IV proxies: (i) 5-minute RV and (ii) duration-based IV estimator derived from the log-ACD-GARCH- s, \tilde{s} model with $K = 78$. Among all models, the HAR model augmented with our duration-based IV estimators (referred to as the “duration-based HAR”) consistently achieves substantially lower MSE and relatively reduced QLIKE results, as shown in Panel A. Since the MSE function heavily penalizes outliers and is highly sensitive to excessively inaccurate forecasts, the results in Panel A suggest that the duration-based HAR models effectively reduce the occurrence of

Table 10: Daily out-of-sample HAR volatility forecasts

	RV				log-ACD-GARCH- s, \bar{s}			
	RW		EW		RW		EW	
	MSE	QLIKE	MSE	QLIKE	MSE	QLIKE	MSE	QLIKE
Panel A: Duration-based estimators with adaptive PDS ($K = 78$)								
HAR-ACD	2.68	0.42	2.67	0.38	1.52	0.15	1.52	0.14
HAR-log-ACD	2.66	0.42	2.65	0.39	1.51	0.15	1.50	0.15
HAR-log-ACD-GARCH	2.65	0.41	2.64	0.38	1.50	0.15	1.50	0.14
HAR-log-ACD-GARCH- s, \bar{s}	2.66	0.41	2.65	0.38	1.50	0.15	1.50	0.14
Panel B: Return-based estimators with CTS (5-minute)								
HAR-RV	2.96	0.44	2.95	0.41	1.86	0.19	1.86	0.19
HAR-BV	2.99	0.44	2.99	0.41	1.88	0.19	1.88	0.18
HAR-MinRV	3.01	0.46	3.01	0.44	1.88	0.20	1.89	0.19
HAR-MedRV	2.96	0.43	2.96	0.41	1.86	0.18	1.86	0.18
HAR-TRV	2.87	0.42	2.87	0.40	1.75	0.17	1.76	0.16
HARQ-RV	3.07	0.38	3.05	0.36	1.93	0.18	1.91	0.18
SHAR-RV	2.80	0.42	2.75	0.41	1.76	0.18	1.73	0.17
Panel C: Return-based estimators with TTS ($K = 78$)								
HAR-RV	2.81	0.42	2.80	0.39	1.65	0.17	1.65	0.16
HAR-BV	2.92	0.44	2.91	0.40	1.74	0.17	1.73	0.16
HAR-MinRV	3.01	0.45	3.00	0.40	1.82	0.18	1.82	0.17
HAR-MedRV	2.88	0.42	2.86	0.39	1.73	0.17	1.72	0.16
HAR-TRV	2.79	0.43	2.78	0.40	1.63	0.17	1.63	0.16
HARQ-RV	3.40	0.37	3.31	0.34	1.91	0.16	1.86	0.16
SHAR-RV	2.73	0.41	2.72	0.38	1.60	0.16	1.60	0.15
Panel D: Noise-corrected return-based estimators with all tick-level data								
HAR-PRV	2.79	0.40	2.77	0.36	1.63	0.15	1.62	0.14
HAR-PBV	2.82	0.40	2.80	0.37	1.64	0.15	1.63	0.14
HAR-PTRV	2.82	0.40	2.80	0.37	1.64	0.15	1.63	0.14

MSE ($\times 10^5$) and QLIKE of daily out-of-sample volatility forecasts for the SPDR S&P 500 ETF Trust (SPY). The HAR model is re-estimated via OLS with both rolling windows and expanding windows, respectively. Durations are obtained with the K -adaptive thresholds for $K = 78$. Return-based estimators are constructed from returns sampled with both CTS and TTS with the same sampling frequency. For PRV and PBV, the pre-averaging window is set to $\lceil \theta \sqrt{n} \rceil$ with $\theta = 0.5$.

extremely misinformative forecasts in the tails.

In Panels B and C, which evaluate standard HAR models with return-based estimators, the HAR-TRV model achieves the lowest MSEs across all cases. For the HAR-RV extensions, both the HARQ and SHAR models deliver improved QLIKE results compared to the original HAR-RV model, while the HARQ model tends to produce more inaccurate forecasts in the tails, as indicated by its higher MSE results. A comparison between Panels B and C reveals that HAR models augmented with TTS-based realized estimators outperform those with CTS-based estimators. This improvement highlights the advantages of incorporating trade intensity adapted to the volatility level, which helps mitigate the underperformance of return-based estimators, aligning with the prior findings by Andersen et al. (2012) and Hautsch and Podolskij (2013). Finally, as shown in Panel D, the HAR-PRV and HAR-PBV models, which utilize all tick-level transaction data, offer substantial improvements over both CTS- and TTS-based methods, which achieve performance comparable to the duration-based HAR models.

Table 11 reports the p -values of the modified Diebold-Mariano test based on MSE, which compares

Table 11: Diebold-Mariano p -values for HAR-log-ACD-GARCH- s, \tilde{s} volatility forecasts

	RV		log-ACD-GARCH- s, \tilde{s}	
	RW	EW	RW	EW
Panel A: Return-based estimators with CTS (5-minute)				
HAR-RV	0.0627	0.0546	0.0076	0.0065
HAR-BV	0.0237	0.0200	0.0038	0.0031
HAR-MinRV	0.0159	0.0132	0.0025	0.0019
HAR-MedRV	0.0320	0.0271	0.0073	0.0058
HAR-TRV	0.0794	0.0685	0.0128	0.0107
HARQ-RV	0.0081	0.0080	0.0020	0.0013
SHAR-RV	0.1569	0.1842	0.0468	0.0649
Panel B: Return-based estimators with TTS ($K = 78$)				
HAR-RV	0.1710	0.1630	0.1118	0.1099
HAR-BV	0.0357	0.0324	0.0276	0.0275
HAR-MinRV	0.0201	0.0180	0.0155	0.0129
HAR-MedRV	0.1060	0.1104	0.0525	0.0560
HAR-TRV	0.1290	0.1254	0.1125	0.1069
HARQ-RV	0.0063	0.0083	0.0137	0.0184
SHAR-RV	0.2724	0.2465	0.2788	0.2932
Panel C: Noise-corrected return-based estimators with all tick-level data				
HAR-PRV	0.2744	0.3057	0.1165	0.1345
HAR-PBV	0.1331	0.1500	0.0647	0.0747

Diebold-Mariano p -values for HAR-log-ACD-GARCH- s, \tilde{s} volatility forecasts. The null hypothesis is that the accuracy of duration-based IV forecasts derived from the HAR-log-ACD-GARCH- s, \tilde{s} model is inferior to the forecasts from an alternative return-based HAR model.

the predictive accuracy of the duration-based HAR model (based on the log-ACD-GARCH- s, \tilde{s} model) with alternative return-based models. Across all panels in Table 11, most p -values are well below 50%, which indicates a consistent preference for the duration-based HAR forecasts. In Panel A, the p -values for most HAR models with CTS-based realized estimators are below conventional significance levels, particularly for HAR-MinRV, HAR-BV, and HARQ-RV. Similar trends are observed for HAR models with TTS-based estimators in Panel B, although the p -values are generally higher compared to their CTS-based counterparts. In Panel C, the noise-corrected PRV and PBV estimators further reduce the gap between return-based and duration-based HAR forecasts, while they do not surpass the duration-based HAR model in terms of forecast accuracy. Overall, the findings highlight the superior forecasting performance of the duration-based HAR model, particularly compared to return-based models that are less adaptive to intraday volatility dynamics.

To further investigate the impact of price jumps on both duration-based and return-based HAR models, we partition the out-of-sample period into two complementary subsets of days: those with jumps and those without. The presence of jumps is identified with the statistical test proposed by Aït-Sahalia et al. (2012). Table 12 reports the MSE results for all selected HAR models on each subset. Across all models, MSEs are consistently higher on days with jumps, while the duration-based HAR models, particularly the one based on the log-ACD-GARCH- s, \tilde{s} model, are less affected and consistently outperform all return-based models in the presence of jumps.

Motivated by the results in Table 12, we extend the standard HAR model by incorporating

Table 12: MSE results of out-of-sample HAR volatility forecasts on days with/without jumps

Jumps?	RV				log-ACD-GARCH- s, \tilde{s}			
	RW		EW		RW		EW	
	Yes	No	Yes	No	Yes	No	Yes	No
Panel A: Duration-based estimators with adaptive PDS ($K = 78$)								
HAR-ACD	3.13	2.53	3.11	2.53	1.92	1.39	1.91	1.40
HAR-log-ACD	3.10	2.51	3.08	2.52	1.89	1.39	1.88	1.39
HAR-log-ACD-GARCH	3.05	2.52	3.02	2.52	1.88	1.37	1.87	1.38
HAR-log-ACD-GARCH- s, \tilde{s}	2.98	2.54	2.95	2.54	1.82	1.39	1.81	1.40
Panel B: Return-based estimators with CTS (5-minute)								
HAR-RV	3.47	2.78	3.43	2.79	2.28	1.72	2.25	1.73
HAR-BV	3.37	2.87	3.33	2.88	2.20	1.77	2.17	1.78
HAR-MinRV	3.35	2.89	3.33	2.90	2.18	1.78	2.16	1.80
HAR-MedRV	3.41	2.82	3.38	2.82	2.25	1.73	2.21	1.74
HAR-TRV	3.21	2.76	3.19	2.77	2.16	1.62	2.14	1.63
HARQ-RV	3.49	3.03	3.44	2.92	2.43	1.76	2.38	1.76
SHAR-RV	3.21	2.69	3.12	2.63	2.15	1.63	2.09	1.61
Panel C: Return-based estimators with TTS ($K = 78$)								
HAR-RV	3.67	2.52	3.63	2.52	2.40	1.41	2.37	1.41
HAR-BV	3.83	2.62	3.80	2.61	2.49	1.49	2.46	1.49
HAR-MinRV	3.91	2.71	3.86	2.71	2.55	1.58	2.51	1.58
HAR-MedRV	3.73	2.59	3.69	2.59	2.47	1.48	2.43	1.48
HAR-TRV	3.53	2.54	3.50	2.54	2.28	1.42	2.26	1.42
HARQ-RV	4.74	2.96	4.56	2.90	3.06	1.53	2.92	1.51
SHAR-RV	3.23	2.44	3.19	2.43	2.14	1.31	2.10	1.30
Panel D: Noise-corrected return-based estimators with all tick-level data								
HAR-PRV	3.11	2.68	3.08	2.66	2.05	1.49	2.02	1.49
HAR-PBV	3.10	2.73	3.07	2.71	2.02	1.52	1.99	1.51

MSE ($\times 10^5$) of daily out-of-sample volatility forecasts on days with and without jumps for the SPDR S&P 500 ETF Trust (SPY). The HAR model is re-estimated via OLS with both rolling windows and expanding windows, respectively. Durations are obtained with the K -adaptive thresholds for $K = 78$. Return-based estimators are constructed from returns sampled with both CTS and TTS with the same sampling frequency. For PRV and PBV, the pre-averaging window is set to $[\theta\sqrt{n}]$ with $\theta = 0.5$. Jumps are identified with the statistical test of [Aït-Sahalia et al. \(2012\)](#).

both the continuous and jump components of total variation on the RHS, i.e., the HAR-J model of [Andersen et al. \(2007a\)](#). Specifically, for all standard HAR models with jump-robust IV estimators, referred to as HAR-CV models, we extend them to HAR-J-CV models by including $J_d^2 = \max\{RV_d - CV_d, 0\}$ on the RHS to explicitly account for jump variation. Furthermore, [Patton and Sheppard \(2015\)](#) demonstrate that it is beneficial to include signed jump variation components:

$$\Delta J_d^{2+} = (RS_d^+ - RS_d^-) \mathbb{1}_{\{RS_d^+ > RS_d^-\}} \quad \text{and} \quad \Delta J_d^{2-} = (RS_d^+ - RS_d^-) \mathbb{1}_{\{RS_d^+ < RS_d^-\}}, \quad (42)$$

which is denoted as the HAR-SJ model. [Table 13](#) presents the MSE and QLIKE results for all selected HAR-J and HAR-SJ models. The results indicate that incorporating (signed) jump variation significantly reduces both MSE and QLIKE for both duration-based and return-based HAR models, with duration-based models showing more pronounced benefits. Among the HAR-J models, the HAR-J-log-ACD-GARCH- s, \tilde{s} achieves the lowest MSE values. Similarly, within the HAR-SJ models, the one based on the log-ACD-GARCH- s, \tilde{s} model outperforms all others with both the lowest MSE

and QLIKE results across all cases.

Table 13: Daily out-of-sample HAR-J volatility forecasts

	RV				log-ACD-GARCH- s, \bar{s}			
	RW		EW		RW		EW	
	MSE	QLIKE	MSE	QLIKE	MSE	QLIKE	MSE	QLIKE
Panel A: Duration-based estimators with adaptive PDS ($K = 78$)								
HAR-J-ACD	2.07	0.27	2.08	0.26	1.42	0.14	1.42	0.13
HAR-J-log-ACD	2.08	0.27	2.08	0.26	1.42	0.13	1.42	0.13
HAR-J-log-ACD-GARCH	2.02	0.26	2.02	0.25	1.38	0.13	1.39	0.12
HAR-J-log-ACD-GARCH- s, \bar{s}	2.01	0.26	2.02	0.25	1.38	0.13	1.38	0.12
HAR-SJ-ACD	1.72	0.23	1.73	0.22	1.16	0.11	1.17	0.10
HAR-SJ-log-ACD	1.73	0.22	1.73	0.21	1.17	0.10	1.17	0.10
HAR-SJ-log-ACD-GARCH	1.71	0.22	1.71	0.21	1.14	0.10	1.15	0.10
HAR-SJ-log-ACD-GARCH- s, \bar{s}	1.70	0.22	1.71	0.21	1.14	0.10	1.15	0.10
Panel B: Return-based estimators with CTS (5-minute)								
HAR-J-BV	2.45	0.31	2.46	0.29	1.80	0.17	1.81	0.16
HAR-J-MinRV	2.47	0.32	2.48	0.30	1.81	0.17	1.83	0.17
HAR-J-MedRV	2.45	0.30	2.46	0.29	1.79	0.17	1.80	0.16
HAR-J-TRV	2.30	0.27	2.31	0.26	1.66	0.15	1.67	0.15
HAR-SJ-BV	1.92	0.25	1.94	0.24	1.43	0.13	1.45	0.13
HAR-SJ-MinRV	1.92	0.26	1.93	0.24	1.45	0.14	1.47	0.14
HAR-SJ-MedRV	1.89	0.25	1.94	0.25	1.43	0.13	1.44	0.13
HAR-SJ-TRV	1.84	0.23	1.90	0.24	1.35	0.12	1.36	0.12
Panel C: Return-based estimators with TTS ($K = 78$)								
HAR-J-BV	2.23	0.28	2.24	0.27	1.59	0.15	1.59	0.14
HAR-J-MinRV	2.30	0.29	2.31	0.28	1.67	0.15	1.67	0.15
HAR-J-MedRV	2.26	0.29	2.26	0.27	1.61	0.15	1.62	0.15
HAR-J-TRV	2.18	0.29	2.18	0.27	1.52	0.15	1.53	0.14
HAR-SJ-BV	1.90	0.26	1.90	0.25	1.31	0.12	1.32	0.12
HAR-SJ-MinRV	1.99	0.27	1.99	0.26	1.39	0.13	1.39	0.13
HAR-SJ-MedRV	1.91	0.26	1.91	0.25	1.34	0.13	1.34	0.13
HAR-SJ-TRV	1.81	0.25	1.82	0.24	1.24	0.12	1.25	0.12
Panel D: Noise-corrected return-based estimators with all tick-level data								
HAR-J-PBV	2.18	0.25	2.18	0.25	1.48	0.14	1.48	0.14
HAR-SJ-PBV	1.72	0.22	1.73	0.22	1.19	0.11	1.19	0.11

MSE ($\times 10^5$) and QLIKE of daily out-of-sample volatility forecasts for the SPDR S&P 500 ETF Trust (SPY). Both HAR-J and HAR-SJ models, which follow [Andersen et al. \(2007a\)](#) and [Patton and Sheppard \(2015\)](#), are re-estimated via OLS with both rolling windows and expanding windows, respectively. Durations are obtained with the K -adaptive thresholds for $K = 78$. Return-based estimators are constructed from returns sampled with both CTS and TTS with the same sampling frequency. For PRV and PBV, the pre-averaging window is set to $\lceil \theta \sqrt{n} \rceil$ with $\theta = 0.5$.

4.3 Intraday Volatility Dynamics Around FOMC Announcements

The short-term impact of macroeconomic news announcements on high-frequency intraday price, volume and volatility dynamics has received a lot of attention from the financial economics and econometrics literature ([Andersen et al., 2003b, 2007b](#); [Lee and Mykland, 2008](#); [Lee, 2012](#); [Bollerslev et al., 2018, 2021](#)). Some recent macroeconomics literature also identifies monetary shocks based on the assumption that the market volatility tends to spike during specific public news announcements such as those associated with FOMC meetings ([Nakamura and Steinsson, 2018](#)). Here we apply

our duration-based spot volatility estimator to examine the short-term impact of FOMC news announcements on intraday volatility dynamics.

We split our sample period from 2017 to 2022 (with the burn-in period before 2017 excluded) into two subsets: 49 days with pre-scheduled FOMC announcements (“FOMC days”), and those without FOMC announcements (“non-FOMC days”). Fig. 7 illustrates the average annualized spot volatility estimates for 5-minute intervals on both FOMC and non-FOMC days, respectively, calculated based on the log-ACD-GARCH- s, \tilde{s} model with $K = 78$.¹⁰

The results reveal an immediate and pronounced volatility spike in the 5-minute interval starting at 14:00 ET on FOMC days, but it is absent on non-FOMC days. Moreover, post-announcement volatility remains significantly elevated on FOMC days compared to non-FOMC days. These empirical findings are in line with [Bollerslev et al. \(2021, 2024\)](#).

In addition, the estimated spot volatilities across all 5-minute intervals exhibit a clear asymmetric U-shaped or reverse J-shaped pattern over trading hours. This pattern aligns with our simulation specifications in Section 3 and corroborates empirical findings in the literature, e.g., [Christensen et al. \(2018\)](#), and [Andersen et al. \(2018, 2019, 2024\)](#).

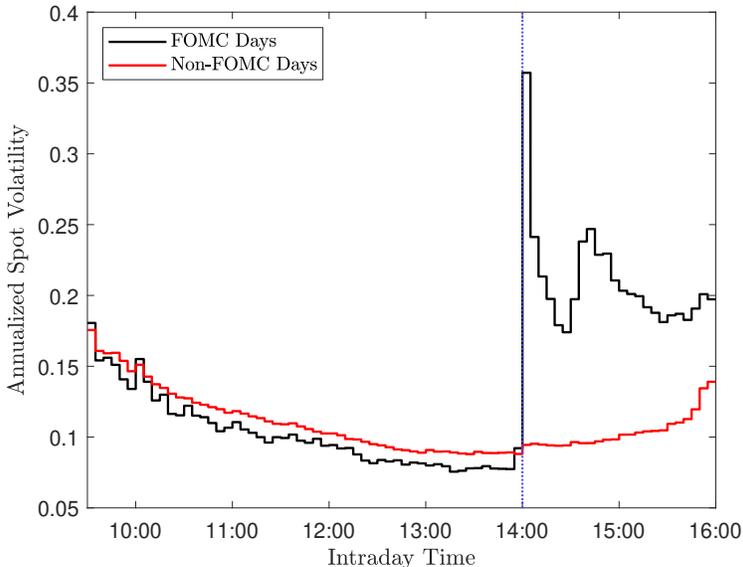


Figure 7: Annualized spot volatility estimates based on the log-ACD-GARCH- s, \tilde{s} model for each equidistant intervals with $\Delta = 5$ minutes. Durations are obtained with the K -adaptive thresholds with $K = 78$.

5 Conclusions

This paper introduces an innovative framework for duration-based volatility estimation. We employ daily adaptive thresholds to separate daily volatility dynamics from the intraday durations over an extended period spanning multiple days. Both spot and integrated volatility estimators are formulated based on classical parametric duration models. This methodology to disentangle daily and intraday volatility dynamics greatly enhances the data availability for model estimation, which

overcomes the limitations of sparse sampling widely adopted in both the financial econometrics literature and by practitioners, and potentially leads to more accurate volatility estimates. The simulation results demonstrate the superior finite-sample performance of our duration-based volatility estimators compared to selected competitors. In an empirical illustration of daily volatility forecasting, we find that the HAR model augmented with our duration-based IV estimator effectively reduces the occurrence of highly misleading forecasts and improves forecasting accuracy according to standard out-of-sample loss functions. Furthermore, based on our spot volatility estimator, we identify an immediate and substantial short-term impact of FOMC news announcements on intraday volatility.

Notes

¹Note that Proposition 1 does not hold under \mathcal{F}_t in the sense that Z_i is not i.i.d. when conditioned on \mathcal{F}_t for $t \in (\tau_{d,i-1}, \tau_{d,i})$, as knowing the value of X_t after $\tau_{d,i-1}$ immediately alters the conditional density of Z_i . However, this issue does not arise under $\overline{\mathcal{F}}_t$, since it does not contain the value of X_t between price events.

²This multiplicative decomposition holds exactly when the price process follows a pure-jump process such as a compounded Poisson process as in [Oomen \(2006\)](#).

³In this section, we also examine three additional duration models within our volatility estimation framework: the ACD model (with Exp(1) residuals), the log-ACD model, and the log-ACD-GARCH model without seasonality effects. The parameter estimates for these models are reported in Appendix B.1. Additional simulation results, including those with different distributions of $u_{d,i}$, various choices of ACD and GARCH lags, and different values of $Q \geq 1$ for the seasonality terms, which do not alter the qualitative results in this section, are available upon request.

⁴Alternatively, the “true” σ_t sampled at the mid-point of each interval can be used, which has minimal impact on the results presented in the subsequent tables and does not affect our conclusions.

⁵Furthermore, we also evaluate other benchmarks, including localized versions of the MinRV and MedRV estimators of [Andersen et al. \(2012\)](#), as well as the differenced-return volatility (DV) estimator of [Andersen et al. \(2023\)](#). The results for these additional benchmarks are provided in Appendix B.2.

⁶The simulation of market microstructure noise can be more realistically calibrated to capture several well-documented empirical features in financial markets, including temporal heteroskedasticity, persistent serial correlation, intraday seasonality, dependence on the latent efficient price, and rounding errors due to discrete tick sizes. In this paper, we focus on the widely studied case of i.i.d. additive Gaussian noise and leave more comprehensive calibration for future research.

⁷We use the SAS code from [Holden and Jacobsen \(2014\)](#) to extract all tick-by-tick transaction records matched with relevant ask/bid quotes from the daily TAQ dataset available on the Wharton Research Data Services (WRDS).

⁸Additional estimation results with various choices of ACD and GARCH lags, and $Q \geq 1$ for the seasonality terms, which do not alter the qualitative results throughout this section, are available upon request.

⁹Following [Bollerslev et al. \(2016\)](#), the “insanity filter” of [Swanson and White \(1997\)](#) is applied: For each rolling or expanding window, the minimum, maximum, and average of in-sample estimates are re-calculated. All one-step-ahead out-of-sample forecasts that are greater (smaller) than the maximum (minimum) in-sample value will be replaced by the in-sample mean.

¹⁰The spot volatility estimates based on alternative duration models are provided in Appendix B.3.

References

- Aït-Sahalia, Y. and Jacod, J. (2009). Testing for jumps in a discretely observed process. *Annals of Statistics*, 37(1):184–222.
- Aït-Sahalia, Y. and Jacod, J. (2014). *High-Frequency Financial Econometrics*. Princeton University Press.
- Aït-Sahalia, Y., Jacod, J., and Li, J. (2012). Testing for jumps in noisy high frequency data. *Journal of Econometrics*, 168(2):207–222.
- Aït-Sahalia, Y. and Kimmel, R. (2007). Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics*, 83(2):413–452.
- Aït-Sahalia, Y., Mykland, P. A., and Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies*, 18(2):351–416.
- Allen, D., Chan, F., McAleer, M., and Peiris, S. (2008). Finite sample properties of the QMLE for the Log-ACD model: Application to Australian stocks. *Journal of Econometrics*, 147(1):163–185.
- Amemiya, T. (1985). *Advanced Econometrics*. Harvard University Press.
- Andersen, T. G. and Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2-3):115–158.
- Andersen, T. G. and Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4):885–905.
- Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007a). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics*, 89(4):701–720.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003a). Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Vega, C. (2003b). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, 93(1):38–62.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Vega, C. (2007b). Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics*, 73(2):251–277.
- Andersen, T. G., Bondarenko, O., Kyle, A. S., and Obizhaeva, A. A. (2018). Intraday trading invariance in the E-mini S&P 500 futures market. Working Paper.

- Andersen, T. G., Dobrev, D., and Schaumburg, E. (2008). Duration-based volatility estimation. Working Paper.
- Andersen, T. G., Dobrev, D., and Schaumburg, E. (2012). Jump-robust volatility estimation using nearest neighbor truncation. *Journal of Econometrics*, 169(1):75–93.
- Andersen, T. G., Li, Y., Todorov, V., and Zhou, B. (2023). Volatility measurement with pockets of extreme return persistence. *Journal of Econometrics*, 237(2):105439.
- Andersen, T. G., Su, T., Todorov, V., and Zhang, Z. (2024). Intraday periodic volatility curves. *Journal of the American Statistical Association*, 119(546):1181–1191.
- Andersen, T. G., Thyrgaard, M., and Todorov, V. (2019). Time-varying periodicity in intraday volatility. *Journal of the American Statistical Association*, 114(528):1695–1707.
- Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2009). Realized kernels in practice: Trades and quotes. *Econometrics Journal*, 12(3):1–32.
- Barndorff-Nielsen, O. E., Kinnebrock, S., and Shephard, N. (2010). Measuring downside risk – realised semivariance. in T. Bollerslev, J. Russell, and M. Watson, eds., *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, New York: Oxford University Press, 2010.
- Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(2):253–280.
- Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2(1):1–37.
- Barndorff-Nielsen, O. E. and Shiryaev, A. N. (2015). *Change of Time and Change of Measure*. World Scientific Publishing Company.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.
- Bollerslev, T., Li, J., and Li, Q. (2024). Optimal nonparametric range-based volatility estimation. *Journal of Econometrics*, 238(1):105548.
- Bollerslev, T., Li, J., and Liao, Z. (2021). Fixed- k inference for volatility. *Quantitative Economics*, 12(4):1053–1084.
- Bollerslev, T., Li, J., and Xue, Y. (2018). Volume, volatility, and public news announcements. *Review of Economic Studies*, 85(4):2005–2041.
- Bollerslev, T., Patton, A. J., and Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18.

- Borodin, A. N. and Salminen, P. (2002). *Handbook of Brownian Motion - Facts and Formulae*. Springer Basel AG, 2nd Edition.
- Christensen, K., Hounyo, U., and Podolskij, M. (2018). Is the diurnal pattern sufficient to explain intraday variation in volatility? A nonparametric assessment. *Journal of Econometrics*, 205(2):336–362.
- Christensen, K., Oomen, R., and Renò, R. (2022). The drift burst hypothesis. *Journal of Econometrics*, 227(2):461–497.
- Christensen, K., Timmermann, A., and Veliyev, B. (2025). Warp speed price moves: Jumps after earnings announcements. *Journal of Financial Economics*, forthcoming.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196.
- Daley, D. J. and Vere-Jones, D. (2003). *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods*. Springer, 2nd Edition.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007.
- Engle, R. F. and Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66(5):1127–1162.
- Foster, D. P. and Nelson, D. B. (1996). Continuous record asymptotics for rolling sample variance estimators. *Econometrica*, 64(1):139.
- Fukasawa, M. and Rosenbaum, M. (2012). Central limit theorems for realized volatility under hitting times of an irregular grid. *Stochastic Processes and their Applications*, 122(12):3901–3920.
- Gatheral, J., Jaisson, T., and Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, 18(6):933–949.
- Gerhard, F. and Hautsch, N. (2002). Volatility estimation on the basis of price intensities. *Journal of Empirical Finance*, 9(1):57–89.
- Harris, L. (1986). A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics*, 16(1):99–117.
- Hasbrouck, J. (1999). The dynamics of discrete bid and ask quotes. *Journal of Finance*, 54(6):2109–2142.
- Hautsch, N. (2011). *Econometrics of Financial High-Frequency Data*. Springer.

- Hautsch, N. and Podolskij, M. (2013). Preaveraging-based estimation of quadratic variation in the presence of noise and jumps: Theory, implementation, and empirical evidence. *Journal of Business & Economic Statistics*, 31(2):165–183.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343.
- Holden, C. W. and Jacobsen, S. (2014). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *Journal of Finance*, 69(4):1747–1785.
- Hong, S. Y., Nolte, I., Taylor, S. J., and Zhao, X. (2023). Volatility estimation and forecasts based on price durations. *Journal of Financial Econometrics*, 21(1):106–144.
- Jacod, J., Li, Y., Mykland, P. A., Podolskij, M., and Vetter, M. (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and their Applications*, 119(7):2249–2276.
- Koopman, S. J., Lucas, A., and Scharth, M. (2016). Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, 98(1):97–110.
- Lee, S. S. (2012). Jumps and information flow in financial markets. *Review of Financial Studies*, 25(2):439–479.
- Lee, S. S. and Mykland, P. A. (2008). Jumps in financial markets: a new nonparametric test and jump dynamics. *Review of Financial Studies*, 21(6):2535–2563.
- Li, Q., Li, Y., Nolte, I., Nolte, S., and Yu, S. (2025). Testing for jumps in a discretely observed price process with endogenous sampling times. Working Paper.
- Li, Y., Nolte, I., and Nolte, S. (2019). Renewal based volatility estimation. Working Paper.
- Li, Y., Nolte, I., and Nolte, S. (2021). High-frequency volatility modeling: A Markov-switching autoregressive conditional intensity model. *Journal of Economic Dynamics and Control*, 124:104077.
- Liu, L. Y., Patton, A. J., and Sheppard, K. (2015). Does anything beat 5-minute RV? a comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1):293–311.
- Mancini, C. (2009). Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics*, 36(2):270–296.
- Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: the information effect. *Quarterly Journal of Economics*, 133(3):1283–1330.
- Oomen, R. C. A. (2006). Properties of realized variance under alternative sampling schemes. *Journal of Business & Economic Statistics*, 24(2):219–237.

- Patton, A. J. and Sheppard, K. (2015). Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics*, 97(3):683–697.
- Pelletier, D. and Wei, W. (2024). A stochastic price duration model for estimating high-frequency volatility. *Journal of Financial Econometrics*, 22(5):1372–1396.
- Podolskij, M. and Vetter, M. (2009). Bipower-type estimation in a noisy diffusion setting. *Stochastic Processes and their Applications*, 119(9):2803–2831.
- Russell, J. R. (1999). Econometric modeling of multivariate irregularly-spaced high-frequency data. Working Paper.
- Swanson, N. R. and White, H. (1997). Forecasting economic time series using flexible versus fixed specification and linear versus nonlinear econometric models. *International Journal of Forecasting*, 13(4):439–461.
- Taylor, S. J. (1982). Financial returns modelled by the product of two stochastic processes – A study of the daily sugar prices 1961-75. In Anderson, O. D., ed., *Time Series Analysis: Theory and Practice*. Vol. 1, North-Holland, 203–226.
- Taylor, S. J. (1986). *Modelling Financial Time Series*. Chichester: John Wiley.
- Taylor, S. J. (1994). Modeling stochastic volatility: A review and comparative study. *Mathematical Finance*, 4(2):183–204.
- Taylor, S. J. (2005). *Asset Price Dynamics, Volatility, and Prediction*. Princeton University Press.
- Tse, Y.-K. and Yang, T. T. (2012). Estimation of high-frequency volatility: An autoregressive conditional duration approach. *Journal of Business & Economic Statistics*, 30(4):533–545.
- Vetter, M. and Zwingmann, T. (2017). A note on central limit theorems for quadratic variation in case of endogenous observation times. *Electronic Journal of Statistics*, 11(1):963–980.
- Wood, R. A., McInish, T. H., and Ord, J. K. (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance*, 40(3):723–739.

Appendix A Proofs

A.1 Proof of Proposition 1

Proof. As the claimed result holds for every $d = 1, 2, \dots$, it suffices to prove the result for $d = 1$ on the interval $[0, 1]$ with $X_0 = 0$ and some threshold δ adapted to \mathcal{F}_0 , and we shall suppress the notation of d in this proof for brevity. We start from the celebrated Dambis-Dubins-Schwarz theorem, which states that all continuous martingales are time-changed Brownian motions under the IV clock (Barndorff-Nielsen and Shiryaev, 2015), i.e., $X_t = \int_0^t \sigma_s dW_s = \widetilde{W}_{V_t}$ is a standard Brownian motion under the V -time $V_t = \int_0^t \sigma_s^2 ds$. Consider the price durations $(x_i)_{1 \leq i \leq N_d}$ generated by Eq. (3), its V -time counterpart can be generated by the same algorithm under a time change:

$$\Delta_i V = \inf_{s>0} \left\{ \left| \widetilde{W}_{V_{i-1}+s} - \widetilde{W}_{V_{i-1}} \right| \geq \delta \right\}, \quad V_i = \sum_{j=1}^i \Delta_j V, \quad (\text{A.1})$$

since the stopping rule commutes with a time change, i.e., $|\widetilde{W}_{V_i} - \widetilde{W}_{V_{i-1}}| = |X_{\tau_i} - X_{\tau_{i-1}}| = \delta$. By the Brownian scaling law:

$$Z_i = \frac{\Delta_i V}{\delta^2} = \inf_{s>0} \left\{ \left| \widetilde{W}_{\widetilde{V}_{i-1}+s}^* - \widetilde{W}_{\widetilde{V}_{i-1}}^* \right| \geq 1 \right\}, \quad (\text{A.2})$$

where $\widetilde{W}_{\widetilde{V}_t}^* = \delta^{-2} \widetilde{W}_{V_t/\delta^2}$ is again a Brownian motion that normalizes the stopping threshold to 1. By the strong Markov property of the Brownian motion, Z_i is independent of \mathcal{F}_{i-1} and $\delta \in \mathcal{F}_0$. The time homogeneity of Brownian motion implies that Z_i must have the same distribution of the first exit time of a Brownian motion B with respect to a symmetric unit interval, i.e., $\inf_{t>0} \{|B_t| \geq 1\}$, and hence (Z_i) is a sequence of i.i.d. random variables. This completes the proof. \square

A.2 Proof of Proposition 2

Proof. Similar to the proof of Proposition 1, we shall assume $d = 1$, $X_0 = 0$, and suppress the notation of d for brevity. We start with an analysis on the sub- σ -field $\overline{\mathcal{F}}_t$. We clearly have $\overline{\mathcal{F}}_0 = \mathcal{F}_0$, which implies that $\delta \in \overline{\mathcal{F}}_0$. Moreover, $N(t)$ is adapted to $\overline{\mathcal{F}}_t$ by construction. Since by assumption $V(t)$ is adapted to $\overline{\mathcal{F}}_t$, the counting process under the IV clock, $\widetilde{N}(V(t))$, is also adapted to $\overline{\mathcal{F}}_t$. Specifically, both $N(t)$ and $\widetilde{N}(V(t))$ are $\overline{\mathcal{F}}_t$ -submartingales, which implies the following Doob-Meyer decompositions:

$$N(t) = M(t) + \Lambda(t), \quad \widetilde{N}(V(t)) = \widetilde{M}(V(t)) + \widetilde{\Lambda}(V(t)), \quad (\text{A.3})$$

where $M(t) = \widetilde{M}(V(t))$ are $\overline{\mathcal{F}}_t$ -martingales, and $\Lambda(t) = \widetilde{\Lambda}(V(t))$ are the compensators for $N(t)$ and $\widetilde{N}(V(t))$, which are $\overline{\mathcal{F}}_t$ -predictable strictly increasing processes with the following representations:

$$\Lambda(t) = \int_0^t \lambda_s ds, \quad \widetilde{\Lambda}(V(t)) = \int_0^{V(t)} \widetilde{\lambda}_{V(s)} dV(s), \quad (\text{A.4})$$

where λ_t and $\tilde{\lambda}_{V(t)}$ are $\overline{\mathcal{F}}_t$ -adapted, positive-valued, and càglàd processes known as the $\overline{\mathcal{F}}_t$ -conditional intensity processes of $N(t)$ and $\tilde{N}(V(t))$, respectively (Hautsch, 2011). Since the two integrals above are identical for all t , we have for any h :

$$\int_t^{t+h} \lambda_s ds = \int_{V(t)}^{V(t+h)} \tilde{\lambda}_{V(s)} dV(s). \quad (\text{A.5})$$

With the definition of conditional intensity (Definition 4.1, Hautsch, 2011), we have for almost all $t \geq 0$:

$$\begin{aligned} \lambda_t &= \lim_{h \downarrow 0} \frac{1}{h} \int_t^{t+h} \lambda_s ds = \lim_{h \downarrow 0} \frac{1}{h} \int_{V(t)}^{V(t+h)} \tilde{\lambda}_{V(s)} dV(s) \\ &= \lim_{h \downarrow 0} \frac{V(t+h) - V(t)}{h} \lim_{h \downarrow 0} \frac{1}{V(t+h) - V(t)} \int_{V(t)}^{V(t+h)} \tilde{\lambda}_{V(s)} dV(s) \\ &= \sigma_t^2 \tilde{\lambda}_{V(t)}, \end{aligned} \quad (\text{A.6})$$

where we utilize the fact that σ_t^2 is a càdlàg process to deduce the first limit above, and one should verify that both sides of the above equation is adapted to $\overline{\mathcal{F}}_t$ by assumption. In particular, the equality $\lambda_t = \sigma_t^2 \tilde{\lambda}_{V(t)}$ holds for almost all $t \in [0, 1]$ except a set with Lebesgue measure zero, which is enough to ensure the equality of the corresponding integrals.

To connect the conditional intensity process with the conditional density of the durations, we use the relationship between the conditional hazard function and the conditional intensity, see, e.g., Eq. (4.1) in Hautsch (2011). In detail, we can write λ_t and $\tilde{\lambda}_{V(t)}$ in terms of the $\overline{\mathcal{F}}_t$ -conditional hazard function of x_i and $\Delta_i V$, i.e., for all $h \in (0, x_i]$:

$$\lambda_{\tau_{i-1}+h} = \frac{f(h|\overline{\mathcal{F}}_{i-1})}{1 - F(h|\overline{\mathcal{F}}_{i-1})}, \quad \tilde{\lambda}_{V(\tau_{i-1}+h)} = \frac{\tilde{f}(\Delta_i V(h)|\overline{\mathcal{F}}_{i-1})}{1 - \tilde{F}(\Delta_i V(h)|\overline{\mathcal{F}}_{i-1})}, \quad (\text{A.7})$$

where $\tilde{f}(\cdot|\overline{\mathcal{F}}_{i-1})$ and $\tilde{F}(\cdot|\overline{\mathcal{F}}_{i-1})$ are the $\overline{\mathcal{F}}_{i-1}$ -conditional PDF and CDF of $\Delta_i V$, respectively. Integrating the conditional intensities over $[\tau_{i-1}, \tau_{i-1} + h]$, we find that:

$$\begin{aligned} -\ln(1 - F(h|\overline{\mathcal{F}}_{i-1})) &= \int_0^h \frac{f(s|\overline{\mathcal{F}}_{i-1})}{1 - F(s|\overline{\mathcal{F}}_{i-1})} ds \\ &= \int_0^{\Delta_i V(h)} \frac{\tilde{f}(s|\overline{\mathcal{F}}_{i-1})}{1 - \tilde{F}(s|\overline{\mathcal{F}}_{i-1})} ds = -\ln(1 - \tilde{F}(\Delta_i V(h)|\overline{\mathcal{F}}_{i-1})), \end{aligned} \quad (\text{A.8})$$

from which we deduce that $F(h|\overline{\mathcal{F}}_{i-1}) = \tilde{F}(\Delta_i V(h)|\overline{\mathcal{F}}_{i-1})$ for all i and $h \in (0, x_{d,i}]$. Finally, it suffices to notice that, conditioning on $\overline{\mathcal{F}}_{i-1} \subset \mathcal{F}_{i-1}$, it holds that $\Delta_i V \stackrel{\mathcal{L}}{=} \delta^2 Z_i$ where Z_i is independent of $\overline{\mathcal{F}}_{i-1}$. Therefore, it implies that

$$F(h|\overline{\mathcal{F}}_{i-1}) = \tilde{F}(\Delta_i V(h)|\overline{\mathcal{F}}_{i-1}) = F_Z(\delta^{-2} \Delta_i V(h)), \quad (\text{A.9})$$

which proves Eq. (8) with the inverse function of $F_Z(\cdot)$.

For Eq. (9), we utilize the fact that $\lambda_t = \sigma_t^2 \tilde{\lambda}_{V(t)}$ holds for almost all t (see Eq. (A.6)). Substituting Eq. (A.7) into this, with the Jacobian transformation $\delta^{-2} f_Z(\delta^{-2} \Delta_i V(h)) = \tilde{f}(\Delta_i V(h) | \bar{\mathcal{F}}_{i-1})$, we find that, for almost all $h \in (0, x_i]$:

$$\sigma_{\tau_{i-1}+h}^2 = \frac{\lambda_{\tau_{i-1}+h}}{\tilde{\lambda}_{V(\tau_{i-1}+h)}} = \frac{f(h | \bar{\mathcal{F}}_{i-1})}{\tilde{f}(\Delta_i V(h) | \bar{\mathcal{F}}_{i-1})} = \frac{\delta^2 f(h | \bar{\mathcal{F}}_{i-1})}{f_Z(\delta^{-2} \Delta_i V(h))}. \quad (\text{A.10})$$

Substituting $\Delta_i V(h) = \delta^2 F_Z^{-1}(F(h | \bar{\mathcal{F}}_{i-1}))$ into the above yields the desired result. This completes the proof. \square

A.3 Proof of Corollary 1

Proof. Similar to the proof of Proposition 1, we shall assume $d = 1$, $X_0 = 0$, and suppress the notation of d for brevity. The relation $x_i = \delta^2 Z_i / \gamma_i$ implies the following relations of the conditional PDF and CDF of x_i and Z_i :

$$F(h | \bar{\mathcal{F}}_{i-1}) = F_Z \left(\frac{h \gamma_i}{\delta^2} \right), \quad f(h | \bar{\mathcal{F}}_{i-1}) = \frac{\gamma_i}{\delta^2} f_Z \left(\frac{h \gamma_i}{\delta^2} \right). \quad (\text{A.11})$$

We substitute the above $F(h | \bar{\mathcal{F}}_{i-1})$ and $f(h | \bar{\mathcal{F}}_{i-1})$ into Eqs. (8) and (9):

$$\Delta_i V(h) = \delta^2 G_i(h) = \delta^2 F_Z^{-1}(F_Z(\delta^{-2} h \gamma_i)) = h \gamma_i, \quad (\text{A.12})$$

$$\sigma_{\tau_{i-1}+h}^2 = \frac{\delta^2 f(h | \bar{\mathcal{F}}_{i-1})}{f_Z(G_i(h))} = \frac{\gamma_{i-1} f_Z(\delta^{-2} h \gamma_i)}{f_Z(F_Z^{-1}(F_Z(\delta^{-2} h \gamma_i)))} = \gamma_i. \quad (\text{A.13})$$

This completes the proof. \square

Appendix B Supplementary Materials

B.1 Estimation Results of Other Duration Models

In addition to the log-ACD-GARCH- s, \tilde{s} model considered in Section 3.2, we also consider three additional duration models: the ACD model (with Exp(1) residuals), the log-ACD model, and the log-ACD-GARCH model without seasonality effects. Table B.1 summarizes the parameter estimates for these models.

Table B.1: Parameter estimates of other models (without seasonality)

Parameters	$K = 78$	$K = 39$	$K = 26$	Parameters	$K = 78$	$K = 39$	$K = 26$
Panel A: ACD(1, 1)							
ϕ_1	0.1299 (0.0056)	0.1595 (0.0102)	0.1697 (0.0146)				
θ_1	0.7996 (0.0085)	0.6682 (0.0205)	0.5285 (0.0368)				
ν_0	21.3672 (1.5878)	105.5405 (9.4895)	281.9996 (29.0604)				
log-likelihood	-156973	-85884	-59351				
Panel B: log-ACD(1, 1)							
ϕ_1	0.9213 (0.0046)	0.8345 (0.0023)	0.7130 (0.0279)				
θ_1	-0.7933 (0.0069)	-0.6773 (0.0083)	-0.5392 (0.0334)				
ν_0	0.4193 (0.0244)	0.9998 (0.0155)	1.8549 (0.1806)				
log-likelihood	-29166	-14581	-9669				
Panel C: log-ACD-GARCH(1, 1, 1, 1)							
ϕ_1	0.9222 (0.0043)	0.8351 (0.0185)	0.7130 (0.0280)	α_1	0.0108 (0.0052)	0.0041 (0.0060)	0.0009 (5.0934)
θ_1	-0.7951 (0.0066)	-0.6786 (0.0182)	-0.5393 (0.0336)	β_1	0.6194 (0.0237)	0.8371 (0.0753)	0.6882 (0.0105)
ν_0	0.4148 (0.0230)	0.9962 (0.1088)	1.8548 (0.1810)	$\tilde{\nu}_0$	0.2587 (0.0170)	0.1142 (0.0503)	0.2325 (3.8048)
log-likelihood	-29164	-14581	-9669				

Parameter estimates (standard errors in parentheses) for three additional duration models: the ACD model (with Exp(1) residuals), the log-ACD model, and the log-ACD-GARCH model without seasonality effects. The DGP follows the Heston model in Eq. (25). Durations are obtained with the K -adaptive thresholds for $K = 78, 39,$ and 26 .

B.2 Monte Carlo RMSE Results of Other Estimators

In addition to the localized return-based estimators in Table 3 and Table 5, we also evaluate the finite-sample performance of the localized versions of MinRV and MedRV proposed by Andersen et al. (2012), as well as the differenced-return volatility (DV) estimator of Andersen et al. (2023). The RMSE results for these return-based estimators are reported in Table B.2 and Table B.3, corresponding to scenarios without and with market microstructure noise, respectively.

Table B.2: Monte Carlo RMSEs for return-based spot volatility estimation

Panel A: All Out-of-Sample Days (300 Days)												
Interval	localized MinRV				localized MedRV				localized DV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	2.15	4.83	11.89	17.02	1.90	4.29	10.79	15.62	2.34	4.37	10.23	14.52
10 min	1.52	3.39	8.43	12.01	1.34	3.00	7.55	10.94	1.95	3.26	7.28	10.36
15 min	1.25	2.76	6.88	9.75	1.10	2.44	6.15	8.78	1.79	2.81	6.00	8.48
30 min	0.89	1.97	4.88	6.88	0.78	1.74	4.31	6.13	1.57	2.22	4.37	6.11
60 min	0.67	1.51	3.73	5.29	0.60	1.33	3.28	4.70	1.33	1.80	3.45	4.74
Panel B: Days with Jumps (51 Days)												
Interval	localized MinRV				localized MedRV				localized DV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	2.24	4.90	12.26	17.51	1.97	4.37	11.10	16.05	2.31	4.49	10.21	14.99
10 min	1.54	3.43	8.67	12.12	1.35	3.10	7.70	11.04	1.88	3.40	7.10	10.70
15 min	1.27	2.83	7.08	9.76	1.12	2.55	6.20	8.87	1.71	2.93	5.90	8.75
30 min	0.87	2.03	4.90	6.72	0.76	1.85	4.16	6.04	1.49	2.29	4.18	6.27
60 min	0.67	1.55	3.85	5.19	0.59	1.39	3.34	4.61	1.27	1.82	3.43	4.88
Panel C: Days without Jumps (249 Days)												
Interval	localized MinRV				localized MedRV				localized DV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	2.13	4.82	11.81	16.92	1.89	4.27	10.73	15.53	2.35	4.34	10.23	14.54
10 min	1.51	3.38	8.38	11.99	1.34	2.98	7.52	10.92	1.96	3.23	7.32	10.28
15 min	1.24	2.75	6.84	9.75	1.10	2.42	6.14	8.76	1.80	2.79	6.02	8.42
30 min	0.89	1.96	4.88	6.91	0.79	1.72	4.34	6.15	1.59	2.20	4.41	6.08
60 min	0.67	1.50	3.71	5.31	0.60	1.32	3.26	4.72	1.34	1.80	3.46	4.71

RMSE results ($\times 10^2$) for the localized return-based estimators based on 1-second, 5-second, 30-second, and 1-minute returns. The DGP follows the Heston model in Eq. (25). The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each equidistant interval.

Table B.3: Monte Carlo RMSEs for return-based spot volatility estimation under market microstructure noise

Panel A: All Out-of-Sample Days (300 Days)												
Interval	localized MinRV				localized MedRV				localized DV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	61.61	20.22	15.41	19.25	58.95	19.51	14.48	18.09	26.75	18.01	13.14	16.15
10 min	61.54	19.77	12.62	14.84	58.90	19.15	12.00	13.94	26.75	17.94	11.25	12.73
15 min	61.52	19.62	11.53	12.96	58.88	19.02	11.06	12.21	26.75	17.91	10.56	11.44
30 min	61.50	19.45	10.36	10.86	58.86	18.88	10.07	10.39	26.70	17.84	9.81	9.95
60 min	61.78	19.49	9.89	9.87	59.12	18.93	9.69	9.60	26.33	17.75	9.51	9.30
Panel B: Days with Jumps (51 Days)												
Interval	localized MinRV				localized MedRV				localized DV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	59.83	18.89	15.11	19.20	57.19	18.22	14.06	17.90	25.67	16.81	12.97	16.19
10 min	59.76	18.41	12.20	14.72	57.13	17.81	11.51	13.69	25.68	16.70	10.85	12.56
15 min	59.74	18.28	10.85	12.68	57.11	17.69	10.30	11.91	25.67	16.65	9.94	11.12
30 min	59.72	18.09	9.67	10.65	57.09	17.54	9.30	10.05	25.62	16.58	9.14	9.50
60 min	60.02	18.09	9.22	9.46	57.38	17.55	8.91	8.99	25.22	16.45	8.70	8.80
Panel C: Days without Jumps (249 Days)												
Interval	localized MinRV				localized MedRV				localized DV			
	1s	5s	30s	1 min	1s	5s	30s	1 min	1s	5s	30s	1 min
5 min	61.96	20.48	15.47	19.26	59.31	19.77	14.53	18.13	26.96	18.25	13.18	16.15
10 min	61.90	20.04	12.70	14.86	59.26	19.41	12.10	13.99	26.97	18.18	11.33	12.77
15 min	61.88	19.88	11.66	13.02	59.24	19.28	11.20	12.27	26.96	18.15	10.69	11.51
30 min	61.86	19.72	10.49	10.90	59.22	19.15	10.22	10.46	26.91	18.09	9.95	10.04
60 min	62.13	19.77	10.02	9.95	59.47	19.20	9.84	9.72	26.55	18.01	9.67	9.40

RMSE results ($\times 10^2$) for the localized return-based estimators based on 1-second, 5-second, 30-second, and 1-minute returns. The RMSE results are calculated based on the annualized spot volatility estimate and the local average of all tick-level spot volatilities (as true value) over each interval.

B.3 Empirical Spot Volatility Estimation with Other Duration Models

Similar to the analysis of intraday volatility around FOMC announcements in Section 4.3, Fig. B.1 presents spot volatility estimates based on three additional duration models: (i) the ACD model, (ii) the log-ACD model, and (iii) the log-ACD-GARCH model (with no seasonality component). All these alternative estimators produce relatively consistent results, which further support the conclusions summarized in Section 4.3. We note that, on FOMC days, the estimates based on the ACD model exhibit a slightly more pronounced spike compared to those from the other three models. This minor discrepancy likely reflects finite-sample variation and does not imply any model deficiency, while it suggests potential scope for refinement in future research.

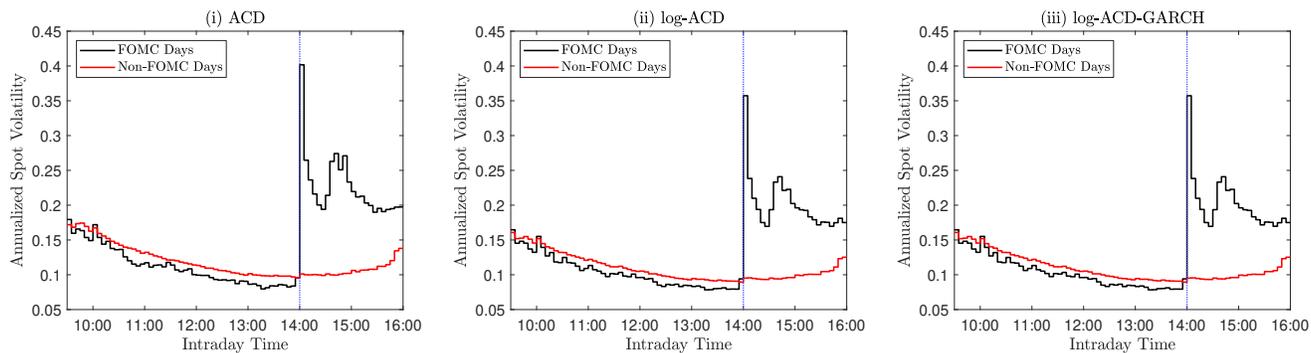


Figure B.1: Annualized spot volatility estimates based on (i) the ACD model, (ii) the log-ACD model, and (iii) the log-ACD-GARCH model (with no seasonality component), for each equidistant intervals with $\Delta = 5$ minutes. Durations are obtained with the K -adaptive thresholds with $K = 78$.