International Journal of
Financial Studies

MDPI

Article

Spot Volatility Measurement Using a Change-Point Duration Model in the High-Frequency Market

Zhicheng Li [1], Haipeng Xing [2] and Yan Wang [3,*]

[1] School of Economics and Trade, Hunan University, Changsha 410079, China; zhicheng.li@hnu.edu.cn
[2] Department of Applied Math, Stony Brook University, Stony Brook, NY 11794, USA; haipeng.xing@stonybrook.edu
[3] Business School, Beijing Normal University, Beijing 100875, China
* Correspondence: wangyan.econ@bnu.edu.cn

Abstract

Modeling high-frequency volatility is an important topic of market microstructure, as it provides the empirical tools to measure and analyze the rapid price movements. Yet, volatility at a high frequency often exhibits abrupt shifts driven by news and trading activity, making accurate estimation challenging. This study develops a change-point duration (CPD) model to estimate spot volatility, in which price-change intensities remain constant between events but may shift at random change points. Using simulations and empirical analysis of Nasdaq limit order book data, we demonstrate that the CPD model achieves a favorable balance between responsiveness to sudden shocks and stability in volatility dynamics. Moreover, it outperforms benchmark approaches, including the classical autoregressive conditional duration model, nonparametric duration-based estimators, and candlestick-based measures. These findings highlight the CPD framework as an effective tool for volatility estimation in high-frequency trading environments.

Keywords: spot volatility; market microstructure; change-point model; price duration

Academic Editor: Vassilios Papavassiliou

Received: 18 August 2025
Revised: 26 September 2025
Accepted: 28 September 2025
Published: 3 October 2025

Citation: Li, Z., Xing, H., & Wang, Y. (2025). Spot Volatility Measurement Using a Change-Point Duration Model in the High-Frequency Market. *International Journal of Financial Studies*, *13*(4), 186. https://doi.org/10.3390/ijfs13040186

1. Introduction

Volatility in modern electronic markets is highly time-varying, often exhibiting sudden spikes and clustering driven by news shocks, algorithmic trading, and shifts in market sentiment. High-frequency volatility modeling is a core topic in market microstructure because it directly captures price dynamics at the most granular level. Accurate high-frequency volatility estimates are essential for market makers, algorithmic traders, and regulators to manage inventory risk, detect regime shifts, and safeguard market stability against rapid price movements and noise. However, many existing estimators remain inadequate, as they often fail to capture both persistence and responsiveness in volatility at fine intervals.

Pioneer studies have focused on the realized volatility estimators to measure integrated variance over fixed intervals (Andersen & Bollerslev, 1998; Andersen et al., 2003; Barndorff-Nielsen & Shephard, 2002, 2004). However, these integrated variance measures smooth out short-lived fluctuations and fail to capture the precise timing of volatility changes. More recently, growing attention has shifted toward spot volatility estimation. Bollerslev et al. (2024) and J. Li et al. (2024) have independently proposed distinct nonparametric estimators for spot volatility, utilizing intraday range data (or candlestick charts).[1] Although the newly proposed nonparametric estimators for spot volatility offer convenient tools for identifying volatility jumps, their time precision remains insufficient for high-frequency traders (HFTs) to effectively adapt to the current trading environment and manage their strategies.

Nowadays, limit order book (LOB) is the dominant mechanism for matching buyers and sellers, valued for its transparency, efficiency, and support of algorithmic trading. It provides rich information about market microstructure as it records millions of quotes and trades with precise time stamps for order submissions, cancellations, and executions. From these time stamps, various forms of duration data can be constructed, such as inter-trade duration (the time interval between consecutive trades), price duration (the time interval that price changes exceed a specified threshold), and others. By exploiting this rich data source, we propose a new spot volatility estimator designed to capture market dynamics at the most granular level.

According to the asymmetric information-based models in market microsturcture theory, "informed traders" trade due to private information on the fundamental value and "liquidity traders" trade due to exogenous reasons (Easley & O'hara, 1992; Glosten & Milgrom, 1985). Traders arrive sequentially to the market, and these arrivals are typically modeled as Poisson processes with different intensities for informed and uninformed participants (Easley et al., 2008, 2002). Consequently, the durations between trades reveal information about the underlying event arrival intensities. A pioneering work for modeling duration data is the autoregressive conditional duration (ACD) model developed by Engle and Russell (1998), which specifies the conditional expectation of duration as a linear function of past durations and past conditional expectations. Subsequently, many extensions of the ACD model were developed to capture nonlinear features of duration data (Bauwens & Giot, 2000; Hujer et al., 2002; Zhang et al., 2001).

Furthermore, if we assume that the underlying price process satisfies Itô semimartingale, then volatility can be derived from the duration between price-change events (price duration). The use of price duration for short-term volatility measurement remains relatively uncommon in the literature. Cho and Frees (1988) were among the first to explore this idea, followed by Gerhard and Hautsch (2002), who carefully introduced a price-duration-based volatility estimator. Then, Tse and Yang (2012) extended this line of research within an ACD model framework. Later, Hong et al. (2023) designed a nonparametric duration-based estimator and found it to be more efficient compared with noise-robust realized volatility estimators.

This paper introduces a change-point price duration (CPD) model for estimating spot volatility using high-frequency price durations, constructed from the LOB data. Change-point models address the common challenge of modeling time series whose parameters may shift occasionally—a concept first explored in the seminal work of Box and Tiao (1975). Such models have broad applications across fields like engineering, econometrics, and biomedicine. A more general statistical framework for the change-point estimation method was further studied by Lai et al. (2005), who introduced the Bounded Complexity Mixture (BCMIX) method to improve computational efficiency. Building on this foundation, Z. Li and Xing (2022) applied the framework to measure quote volatility and to evaluate the liquidity cost. In this study, we extend the approach by modeling jumps in the intensity of price movements as change points and explicitly linking the intensity to spot volatility.

The CPD model assumes an infinite domain for intensity renewal and incorporates a continuous renewal distribution, providing a more flexible and expansive modeling framework than the ACD model. Simulation results demonstrate that the CPD model delivers a stable yet adaptive characterization of volatility by accommodating structural breaks. We compare the volatility estimator by our CPD model with the ACD model, the nonparametric duration estimator by Hong et al. (2023), and the optimal candlestick estimator by J. Li et al. (2024) using real Nasdaq data. The results consistently show that our model achieves superior performance in capturing volatility shift due to sudden market

shocks and clustering effects mentioned in the literature on market microstructure (Admati & Pfleiderer, 1988; Hautsch, 2011).

The paper is organized as follows. Section 2 presents our CPD model and the estimator for spot volatility based on price durations. Section 3 provides the estimation algorithm for the CPD model and a simulation study. In Section 4, we apply our model to the Nasdaq LOB data on 15 April 2013 and compare it with the other three methods for estimating high-frequency volatility. Finally, Section 5 concludes this paper.

2. Model

2.1. Spot Volatility Measurement via Price Duration

Assume that the (log) price process P is an Itô semimartingale defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ and represented as

$$P_t = P_0 + \int_0^t b_s\, ds + \int_0^t \sigma_s\, dW_s + J_t, \tag{1}$$

where b is the drift process, σ is the stochastic volatility process, W is a standard Brownian motion, and J is a pure-jump process driven by a Poisson random measure. Our objective is to estimate the spot volatility σ_t at a specific time point t.

Traditional methods for measuring spot volatility were constructed from localized versions of realized variance (Andersen et al., 2003; Barndorff-Nielsen & Shephard, 2002). Under the standard high-frequency setting, the price process P is observed on a regular grid $\{0, \Delta, 2\Delta, \dots\}$ within a fixed interval $[0, T]$. Let $r_i \equiv P_{i\Delta} - P_{(i-1)\Delta}$ denote the i-th return. For spot volatility estimation, select a bandwidth k and the block-based estimator of the spot variance σ_t^2 is

$$\hat{\sigma}_t^2(k) \equiv \frac{1}{k\Delta} \sum_{j=1}^{k} r_{(i-1)k+j}^2. \tag{2}$$

Asymptotically, as the sampling interval $\Delta_n \to 0$, the conditional spot variance at any time t is expressed as follows:

$$\sigma^2(t) := \lim_{\Delta \downarrow 0} \mathbb{E}\left[\frac{1}{\Delta} (P(t+\Delta) - P(t))^2 \middle| \mathcal{F}_t \right], \tag{3}$$

where $\mathcal{F}_t$ represents the information set available at time t.

As shown in Gerhard and Hautsch (2002), Tse and Yang (2012), and Hong et al. (2023), a duration–based variance estimator can be derived from the relationship between the conditional intensity function of a point process and the corresponding spot variance. Treating the sequence of price changes as a point process, let δ denote the price changing threshold and $\{t_i^\delta\}_{i=1,\dots,n}$ denote the event times when the absolute change in price exceeds δ. Thus, n, the number of events is subject to the chosen δ.

Define the price-change duration (or simply price duartion) between consecutive events as

$$y_i^\delta := t_i^\delta - t_{i-1}^\delta.$$

The conditional variance per unit time over the interval $[t_i^\delta, t_{i+1}^\delta)$ is

$$\sigma^2\left(t_i^\delta\right) = \mathbb{E}\left[\frac{1}{x_{i+1}^\delta} \middle| \mathcal{F}_{t_i^\delta} \right] \cdot \delta^2, \tag{4}$$

where $\delta = |P_{t^{\delta}_{i+1}} - P_{t^{\delta}_{i}}|$. According to the point process theory, the associated counting process $N(t)$ records the number of events up to time t. Its conditional intensity function is

$$\lambda(t; \mathcal{F}_t) = \lim_{\Delta \downarrow 0} \frac{1}{\Delta} \Pr\big[N(t+\Delta) > N(t) \,\big|\, \mathcal{F}_t\big], \tag{5}$$

representing the instantaneous probability of a new event arrival.

For a price-change process with threshold δ, the return variation in an infinitesimal interval Δ is either $\pm\delta$. Hence, the spot variance can be rewritten as follows:

$$\begin{aligned} \sigma^2(t) &= \lim_{\Delta \downarrow 0} \frac{1}{\Delta} \Pr\big[|P(t+\Delta) - P(t)| \geq \delta \,\big|\, \mathcal{F}_t\big] \cdot \delta^2 \\ &= \lambda_{\delta}(t; \mathcal{F}_t) \cdot \delta^2. \end{aligned} \tag{6}$$

Equation (6) shows that measuring spot variance (or volatility by taking its square root) reduces to estimating the event arrival intensity $\lambda_{\delta}(t; \mathcal{F}_t)$ for δ-sized price changes. This formulation underpins the construction of our model.

According to Daley and Vere-Jones (2003), Barndorff-Nielsen and Shiryaev (2015), and Hautsch (2011), the integrated intensity satisfies

$$\Lambda(t_{i-1}, t_i) \equiv \int_{t_{i-1}}^{t_i} \lambda(s)\, ds \overset{i.i.d.}{\sim} \text{Exp}(1), \tag{7}$$

where $t_1, t_2, \ldots, t_n$ denotes the event arrival times of a point process, and $\text{Exp}(1)$ denotes the exponential distribution with rate parameter 1.

2.2. The Change-Point Price Duration Model

To measure the spot volatility in a high-frequency environment, we propose the Change-point Price Duration (CPD) model, inspired by the change-point framework of Box and Tiao (1975) and its computationally efficient extensions by Lai et al. (2005). In general, change-point models are designed to capture structural breaks in model parameters over time—changes that may be driven by shifts in market conditions, liquidity shocks, or new information arrivals. In our setting, the parameter of interest is the price-change intensity, which we interpret as the instantaneous rate at which the price changes by at least a fixed threshold δ. A change point corresponds to a sudden shift in this intensity, reflecting an abrupt change in the underlying volatility regime.

In the CPD model, we assume that the price-change intensity remains constant between two consecutive price-change events, i.e., $\lambda(t) = \lambda(t_{i-1})$ for $t_{i-1} \leq t < t_i$. Thus, we will have $\lambda_i y_i \overset{i.i.d.}{\sim} \text{Exp}(1)$ from Equation (7), where $y_i = t_i - t_{i-1}$ is the price-change duration. Rearranging gives the *mixture-of-exponentials representation*:

$$y_i = \frac{\varepsilon_i}{\lambda_i}, \quad \varepsilon_i \overset{i.i.d.}{\sim} \text{Exp}(1). \tag{8}$$

The dynamics of λ_i is modeled as a discrete-time Markov change-point process with a renewal distribution $G(\cdot)$. At each new event time, the intensity either persists at its previous value with probability $1 - p$ or is redrawn from $G(\cdot)$ with probability p:

$$\lambda_{i+1} = \begin{cases} \lambda_i, & \text{with probability } 1-p, \\ \lambda \overset{i.i.d.}{\sim} G(\cdot), & \text{with probability } p. \end{cases} \tag{9}$$

This specification allows the intensity process to remain constant for extended periods—capturing stable market conditions—interspersed with abrupt shifts to new levels, representing sudden changes in volatility regimes.

For flexibility and tractability, we assume $G(\cdot)$ follows a Gamma distribution with shape parameter α and rate parameter β:

$$G(\lambda) = \text{Gamma}(\lambda; \alpha, \beta) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda}, \tag{10}$$

where $\Gamma(\cdot)$ is the Gamma function. The Gamma distribution is a natural choice here because it is the conjugate prior for the exponential likelihood in (8), which facilitates analytical tractability in estimation.

The economic intuition behind the CPD model is also straightforward. When no change point occurs (probability $1 - p$), the market is in a volatility clustering period, and the intensity of price change is unchanged. When a change point occurs (probability p), the market switches into a new regime due to an exogenous shock such as a burst of trading activity, a sudden shift in order flow, or the arrival of important information. This results in an updated price-change intensity and, consequently, a new level of price volatility.

2.3. Comparison with Other Models

2.3.1. Compare with ACD Model

The ACD model proposed by Engle and Russell (1998) specifies the conditional expectation of durations as a linear function of past durations and past expectations, i.e.,

$$\begin{aligned} y_i &= \varphi_i \varepsilon_i \quad \text{with } \varphi_i = \mathrm{E}\left[y_i \middle| \mathcal{F}_t\right] \\ \varphi_i &= \mu + \sum_{j=0}^{m} \alpha_j y_{i-j} + \sum_{j=0}^{n} \beta_j \varphi_{i-j}, \end{aligned} \tag{11}$$

where the error term enters multiplicatively and all past information influences current durations through the conditional mean. The model, which combines transition analysis with Engle's ARCH framework, is motivated by the empirical observation that financial events tend to cluster in time.

The ACD model can be transformed to an equivalent formulation of the conditional intensity model as in Equation (8), by assuming $\varphi_i \equiv 1/\lambda_i$. Thus, the only difference between the CPD and the ACD model is that they assume a different dynamic structure of the underlying intensity. As the ACD model has a smooth-transition intensity, our CPD model can offer two theoretical advantages. First, it is better suited to ultra-high-frequency data, where the price updates often occur in bursts separated by many irregular pauses. The piecewise-constant structure between change points avoids the need to model complex intra-interval dynamics. Second, by allowing the new intensity levels to be drawn from a continuous distribution, the model can accommodate both small and large jumps in volatility.

2.3.2. Compare with Nonparametric Duration-Based Estimator

Compared with the nonparametric duration-based variance (NPDV) estimator of Hong et al. (2023), the CPD model provides a richer description of volatility dynamics. NPDV is simple and robust, relying only on counting the number of price-change events within a fixed interval to approximate integrated variance. While this makes it computationally efficient and resilient to microstructure noise, it cannot capture fine-grained dynamics of spot volatility or respond sharply to sudden market shifts within an interval. In contrast, the CPD framework models price durations in event time and explicitly allows for structural breaks in the underlying intensity, thereby producing more persistent yet responsive spot

volatility estimates. This makes CPD particularly advantageous in ultra-high-frequency environments where clustering and abrupt shifts in activity are common.

2.3.3. Compare with Candlestick-Based Estimator

The candlestick-based (or range-based) volatility estimator uses intraday OHLC (open, high, low, close) data within fixed clock-time intervals to infer spot volatility (Bollerslev et al., 2024; J. Li et al., 2024). These methods are convenient to implement and often perform well at moderate or coarse horizons, where OHLC values are reliably observed. However, in very short intervals of high-frequency data, the candlestick-based estimator may fail to provide accurate spot measures, since the four distinct prices may not exist or may collapse due to sparse trading. As a result, candlestick estimators tend to smooth away short-lived volatility bursts, or even fail to have values, limiting their usefulness for high-frequency trading applications. By contrast, the CPD model operates directly on tick-level durations, ensuring that volatility dynamics are captured continuously and with greater precision in second-level trading environments.

3. Model Estimation and Simulation

3.1. Estimation Algorithm

In this section, we present the estimation algorithm for our CPD model with specification of (8), (9), and (10). Given the observed price-change durations $\{y_1, y_2, \ldots y_t, \ldots y_n\}$, we want to estimate model parameters. Moreover, the underlying intensities λ_t for $t = 1 \ldots n$ are hidden variables that cannot be observed. The Bayesian structure of this change-point process is shown below in Figure 1. Thus, the joint probability of all observations and hidden states is as follows:

$$P(y_{1n}, \lambda_{1n}) = P(\lambda_1) f(y_1|\lambda_1) \cdot \prod_{t=2}^{n} f(y_t|\lambda_t) P(\lambda_t|\lambda_{t-1}), \tag{12}$$

where $y_{1n} = \{y_1, y_2, \ldots y_t, \ldots y_n\}$, $\lambda_{1n} = \{\lambda_1, \lambda_2, \ldots \lambda_t, \ldots \lambda_n\}$, and $f(y_t|\lambda_t) = \lambda_t e^{-\lambda_t y_t}$.

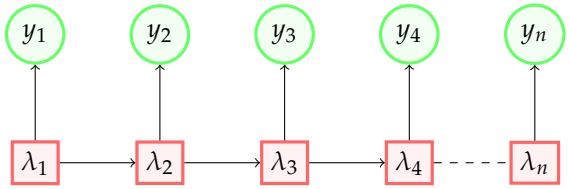

Figure 1. The Bayesian structure of change-point price duration model.

The parameters to be estimated are (α, β, p), corresponding to the shape and rate of the Gamma distribution and the change-point probability, respectively. Since the latent intensity path $\{\lambda_t\}$ is unobserved, direct maximization of the complete log-likelihood is infeasible. We therefore employ the Expectation–Maximization (EM) algorithm for parameter estimation, which is a fully parametric approach.

The complete-data log-likelihood can be written as

$$\ell_c(\{y_t\}, \{\lambda_t\}) = \log G(\lambda_1) + \sum_{t=1}^{n} \log f(y_t \mid \lambda_t) + \sum_{t=2}^{n} \log P(\lambda_t \mid \lambda_{t-1}), \tag{13}$$

where $G(\cdot)$ is the Gamma renewal distribution and $P(\lambda_t \mid \lambda_{t-1})$ represents the probability of λ_t conditional on its previous intensity.

The EM algorithm alternates between the following:

- **E-step:** Compute the expected complete log-likelihood $\mathbb{E}[\ell_c \mid \text{data}]$, given current parameter values. This involves evaluating the posterior change-point probability and the distribution of intensity λ_t using forward–backward filtering (See Appendix A). As the prior distribution $G(\lambda)$ is a Gamma distribution, which is a conjugate prior for the exponential distribution, the posterior distribution of λ_t is also a Gamma distribution. Hence, the posterior distribution of intensity λ_t is as follows:

$$f(\lambda_t|\mathcal{D}) = \sum_{1\le i\le t\le j\le n} \Pi_{itj} \cdot g_{ij}(\lambda_t), \tag{14}$$

  where $g_{ij}(\lambda) \sim \text{Gamma}(\alpha + (j-i+1),\ \beta + \sum_i^j y_t)$ and Π_{itj} represent the change-point probability that the last changing point occurs at i and the next changing point occurs at $j+1$. The calculation steps of this posterior change-point probability is introduced in Appendix B.
- **M-step:** Maximize the expected log-likelihood with respect to α, β, and p. The update for p has a closed form:

$$\hat{p} = \frac{\sum_{t=2}^{n} P(I_t = 1 \mid \text{data})}{n-1}.$$

  $I_t = 1$ indicates there is a changing point at the t-th price-change event, i.e., $\lambda_t \neq \lambda_{t-1}$, and $P(I_t = 1 \mid \text{data})$ represents the probability of a changing point at the t-th price-change event, given the observed data and old parameters.
  Posterior mean estimates of λ_t are then computed as

$$\hat{\lambda}_t = \sum_{i\le t\le j} \Pi_{itj} \frac{\hat{\alpha} + j - i + 1}{\hat{\beta} + \sum_{s=i}^{j} y_s}, \tag{15}$$

  where $(\hat{\alpha}, \hat{\beta})$ are obtained by solving the likelihood equations numerically. The computation of $(\hat{\alpha}, \hat{\beta})$ and Π_{itj} is described in Appendices A and B.

Once the EM algorithm converges, we can again compute the price-change intensities $\{\hat{\lambda}_t^*\}_{t=1}^n$ based on Equation (15) and the converged estimators of $(\hat{\alpha}^*, \hat{\beta}^*)$. Consequently, the estimated spot volatility can be derived from Equation (6). We should note that the volatility is treated to be fixed between consecutive price-change events.

3.2. Simulation Study

We first evaluate the ability of the CPD model to estimate volatility by simulating a price series governed by an underlying stochastic volatility process. The simulated results are shown in Figure 2. As shown, the price exhibits a steady upward drift in the early period, followed by a sudden fall after approximately 600 s. In the meantime, the volatility undergoes a structural shift from a low-volatility period to a high-volatility period. The simulated volatility in Figure 2 is recorded at a sampling frequency of one second.

To apply our model, we construct a series of price-change durations by setting the price-change threshold. Since the minimum tick size of the simulated price is 0.01, we just set the price-change threshold for defining price durations to two ticks, i.e., $\delta = 0.02$. A price-change event is recorded, along with its timestamp, whenever the simulated price moves up or down by 0.02 units. The time interval between consecutive events defines the price duration. Figure 3 presents the resulting price-duration series, together with the corresponding price-change intensities estimated from the CPD model. Since long durations imply low underlying intensities, we plot the reciprocals of the estimated intensities in Figure 3 to facilitate a consistent visual comparison. The results indicate that long (short) durations are associated with low (high) price-change intensities, and that the intensities exhibit notable persistence.

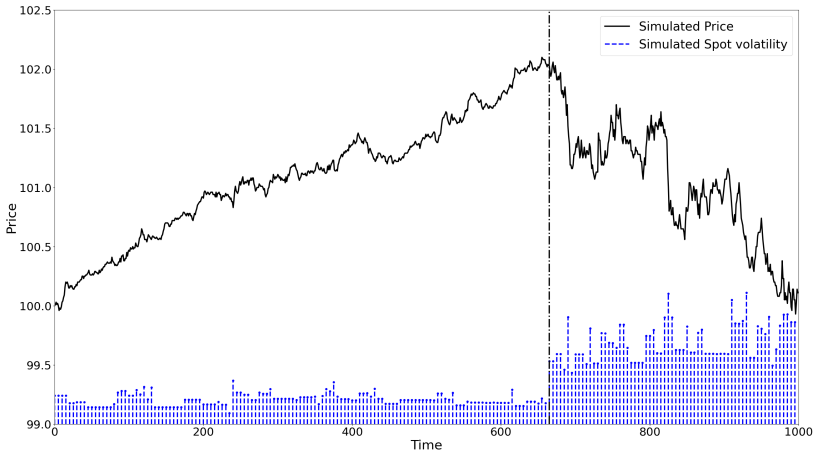

Figure 2. Simulated price and volatility under two regimes, separated by the vertical dashed line.

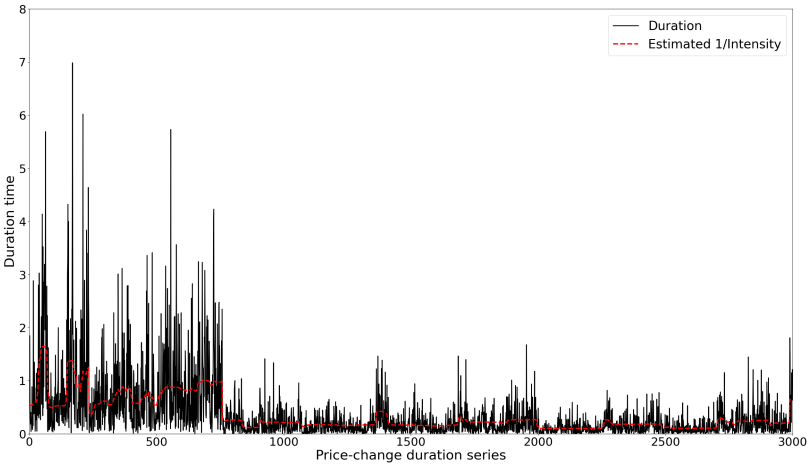

Figure 3. Simulated price durations and estimated intensities.

We then compute the spot volatility using Equation (6) and compare it with the true simulated volatility in Figure 4. From the results, we can find that the estimated volatility closely matches the true values. Moreover, the estimates capture the structural shift from a low-volatility regime to a high-volatility regime.

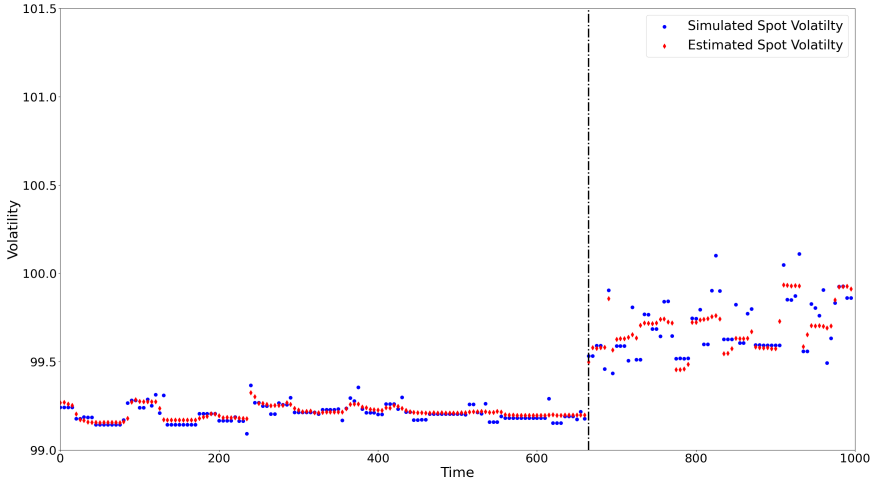

Figure 4. Simulated and estimated volatility under two regimes, separated by the vertical dashed line.

4. Real Data Analysis

4.1. Data Environment

We use Nasdaq LOB data for the real data analysis. The data are obtained from LOBSTER (https://lobsterdata.com/), which provides high-quality LOB data for all Nasdaq stocks beginning in June 2007. The LOB reconstructed by LOBSTER is based on Nasdaq's Historical TotalView-ITCH data. For each active trading day and ticker, LOBSTER produces two files: a 'message' file, which records the events that trigger updates to the LOB within the specified price range, and an 'order book' file, which reports the state of the LOB at each update. We select 15 April 2013 for the one-day analysis because the unexpected Boston Marathon bombing significantly affected the stock market on that trading day, triggering a sharp volatility spike around the event.[2] This setting provides an ideal environment to assess whether our model can capture such abrupt volatility shifts in comparison with other estimators.

Table 1 presents a sample of the message and order book files for Apple Inc. on 15 April 2013. Panel A shows five recorded events at microsecond precision. Event type 1 corresponds to the submission of a limit order, while type 4 indicates the execution of a limit order. The direction indicator '−1' denotes an ask-side event, whereas '1' denotes a bid-side event. Panel B displays the corresponding order book states following each event. For example, at time 34,209.630561 s (third event), the best ask order 10962867 was fully executed with the remaining size of 56 at a price of USD 426.79. After this event, the best ask shifted to the previous second-best level of USD 426.80.

Table 1. The sample message file and order book file of LOBSTER data.

Panel A: Message File					
Time (s)	**Event Type**	**Order ID**	**Size**	**Price**	**Direction**
34,209.630122	1	10962867	100	426.79	−1
34,209.630453	4	10962867	44	426.79	−1
34,209.630561	4	10962867	56	426.79	−1
34,209.630623	1	10962881	100	426.69	1
34,209.633680	1	10962942	100	426.79	−1

Panel B: Order Book File							
Ask Price 1	**Ask Size 1**	**Bid Price 1**	**Bid Size 1**	**Ask Price 2**	**Ask Size 2**	**Bid Price 2**	**Bid Size 2**
426.79	100	426.67	100	426.80	100	426.62	12
426.79	56	426.67	100	426.80	100	426.62	12
426.80	100	426.67	100	426.95	100	426.62	12
426.80	100	426.69	100	426.95	100	426.67	100
426.79	100	426.69	100	426.80	100	426.67	100

Based on these two files, we can construct the price series and the price-change durations. The price is defined as the mid-price, i.e., the average of the best bid and best ask, which can eliminate the bias in duration-based estimators caused by the bid–ask spread that is mentioned in Hong et al. (2023). In practice, the choice of price-change threshold δ involves a trade-off between capturing high-frequency dynamics and mitigating noise. A larger δ produces longer durations and smoother volatility paths, but may overlook short-lived fluctuations. Conversely, a smaller δ improves granularity but may be affected by noise from bid–ask dynamics or order cancellations. As noted by Hong et al. (2023), a threshold of roughly three times the average bid–ask spread balances bias and efficiency in nonparametric estimators. In our CPD framework, bias is less of a concern since we use mid-quote prices to avoid bid–ask effects and construct durations directly from LOB data to avoid time-discretization. Accordingly, we aim to select a δ value small enough to

capture granular price dynamics while relying on the change-point structure of the CPD model to filter out excess noise. For our empirical study, we set $\delta = 0.03$, which strikes this balance and avoids spurious price changes induced by small order cancellations.

Figure 5 plots the series of price durations for Apple stock on 15 April 2013. On that day, there are approximately 9000 durations, and the level of these price duartions shows a very large variation, with the shortest lasting 20 ms and the longest about 70 s.

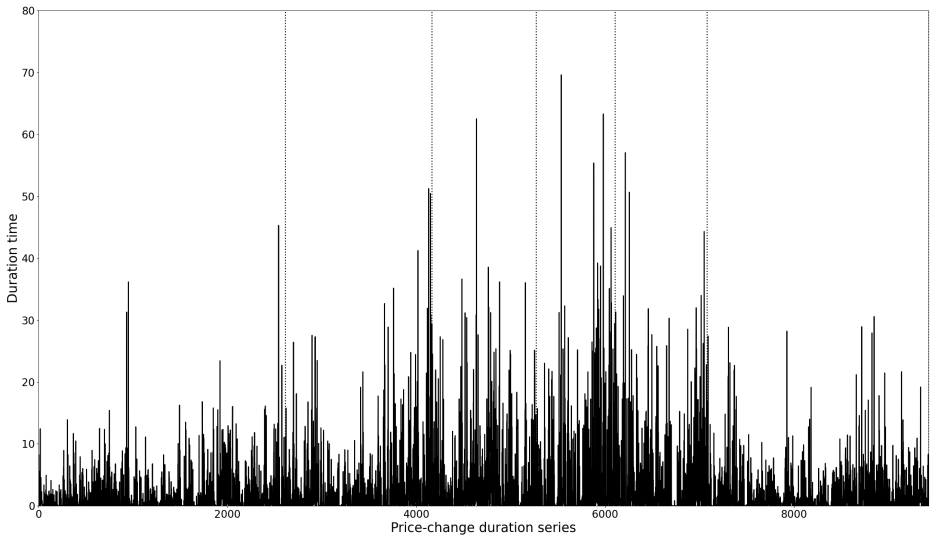

Figure 5. Price duration series.

4.2. Spot Volatility Measurement

We use the constructed price durations as input to the CPD model to estimate spot volatility. According to the estimation algorithm presented in Section 3.1, the parameter estimates are $\hat{\alpha} = 0.23$, $\hat{\beta} = 0.005$, and $\hat{p} = 0.22$. Based on Equation (15), we further infer the latent intensities $\{\lambda_t\}_{1\ldots N}$. The spot volatility $\sigma(t)$ is then derived from Equation (6). Specifically, for $t \in [t_i, t_{i+1})$, we have

$$\sigma^2(t) = \lambda(t_i) \cdot (\delta)^2, \quad t \in [t_i, t_{i+1}),$$

since the model assumes intensity to be constant between two consecutive events.

To evaluate model performance, we compare the CPD model with the classical ACD model introduced in Section 2.3. From the above discussion of the integrated intensity function (Equation (7)) and the mixture of exponential expression (Equation (8)), the estimated residuals $\hat{\varepsilon}_t$ should follow an Exponential(1) distribution. Figure 6 shows that the CPD model residuals adhere closely to this distribution, whereas the ACD residuals deviate substantially. This discrepancy likely arises because the ACD model cannot accommodate the large variability in high-frequency durations, which range from microseconds to tens of seconds.

Moreover, the following Figure 7 further compares the spot volatility estimates from both models over a period of 1200 s, alongside the corresponding price data. The ACD model produces highly volatile estimates with sharp spikes whenever prices change significantly. By contrast, the CPD model generates volatility estimates that display greater persistence while still capturing abrupt price shifts.

In addition to comparing our CPD model with the standard linear ACD model, we also evaluate it against the nonlinear threshold ACD (TACD) model proposed by Zhang et al. (2001). The results are very similar and are presented in Appendix C.

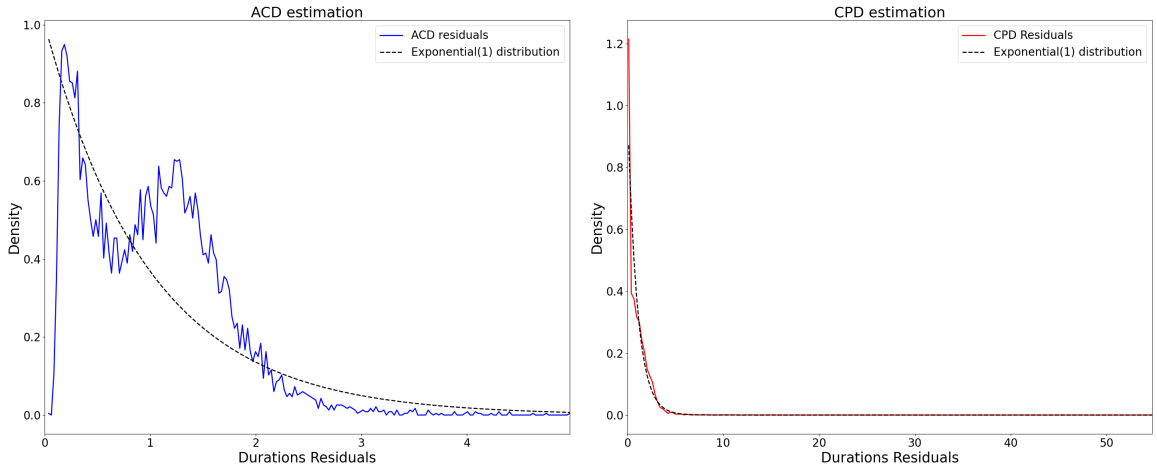

Figure 6. Fitted duration residuals by ACD model and CPD model.

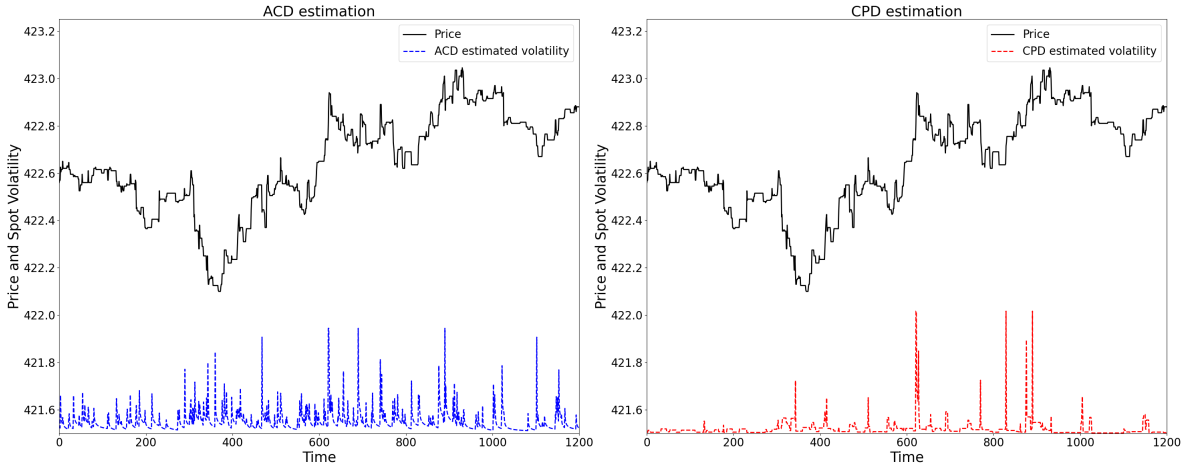

Figure 7. Spot volatility estimation by ACD model and CPD model.

4.3. Integrated Variance Measurement

The integrated variance over the interval $[0, t]$ is as follows:

$$IV(t, t+\Delta t) := \int_t^{t+\Delta t} \sigma^2(\tau) dW_\tau,$$

which is the cumulative variance in returns over a time interval, providing a key measure of total risk. For portfolio managers, it quantifies exposure over the investment horizon and supports risk metrics such as Value-at-Risk. In derivative markets, it underlies option pricing through its link to quadratic variation. For high-frequency market makers, integrated variance guides bid–ask spreads and inventory control, since higher variance signals greater adverse selection risk, while lower variance allows for more aggressive quoting.

In this section, we compare the CPD model with two alternative estimators for high-frequency variance. The first is the nonparametric duration-based variance (NPDV) estimator proposed by Hong et al. (2023). It is defined over the interval $[t, t+\Delta t]$ as

$$NPDV(t, t+\Delta t) = \delta \cdot N(t, t+\Delta t), \tag{16}$$

where δ is the price-change threshold used to define events in both the CPD model and the NPDV, and $N(t, t+\Delta t)$ denotes the number of price-change events occurring in the interval $[t, t+\Delta t]$. As shown in Hong et al. (2023), the NPDV performs comparably to

parametric duration-based estimators such as the ACD model, while outperforming most of the realized variance estimators in high-frequency environments.

The second one is the optimal candlestick estimator by J. Li et al. (2024). This is designed as a nonparametric estimator of 'spot' volatility and given by

$$\frac{0.811 \times (\text{High} - \text{Low}) - |\text{Close} - \text{Open}|}{\text{Duration of Interval}}, \tag{17}$$

where 'High' and 'Low' denote the highest price and lowest price in the interval, and 'Close' and 'Open' are the closing price and opening prices in the interval. As we have mentioned in Section 2.3, this approach is limited in estimating spot volatility at very fine scales, since short intervals may not contain four distinct prices. Consequently, we apply it to construct integrated variance over relatively large intervals.

Together with the CPD and ACD models, we consider four estimators of integrated variance. Figures 8 and 9 present the results for 10 s and 30 s intervals, respectively, over a 1200 s sample period. The CPD estimator produces relatively persistent variance dynamics and effectively captures volatility shifts. By contrast, the ACD estimator generates more volatile patterns with frequent sharp spikes, reflecting sensitivity to abrupt price movements. The NPDV estimator provides a stable fit but tends to understate variance in certain periods, while the candlestick estimator delivers the smoothest series but misses short-lived fluctuations due to its reliance on intraday candlestick prices. Moreover, the nonparametric duration estimator and the candlestick estimator can occasionally take the value of zero, either because no price-change events occur as defined, or because prices remain unchanged within some very short intervals. Therefore, our CPD model performs better by capturing both persistence and responsiveness in high-frequency variance estimation in these short intervals.

Figure 8. Estimated integrated variance over 10 s.

Furthermore, we estimate integrated variances over 1 min and 3 min intervals, with results shown in Figures 10 and 11. To accommodate these longer horizons, the last three trading hours of 15 April 2013 are plotted, with the dashed vertical line marking the Boston Marathon bombing at approximately 3:00 PM, which triggered an immediate market shock on that day. Across all methods, a clear surge in market volatility is observed following the event, confirming the models' sensitivity to sudden volatility jumps. The CPD estimator

responds sharply while maintaining persistence, effectively capturing both the magnitude and duration of the volatility spike. The ACD estimator also detects the shock but produces noisier estimates with more fluctuations. The NPDV estimator exhibits smoother dynamics but appears to underestimate the magnitude of the shock. The candlestick estimator provides a similar patten in the low-volatility regime as our CPD model, while it fails to capture the volatility burst and decaying period following the Boston Marathon bombing. Although the nonparametric NPDV and candlestick estimator are straightforward to compute and attractive for quick, large-scale applications, they also sacrifice flexibility in capturing structural breaks or persistence in volatility. Overall, the CPD model provides the most balanced characterization of market volatility around the event.

Figure 9. Estimated integrated variance over 30 s.

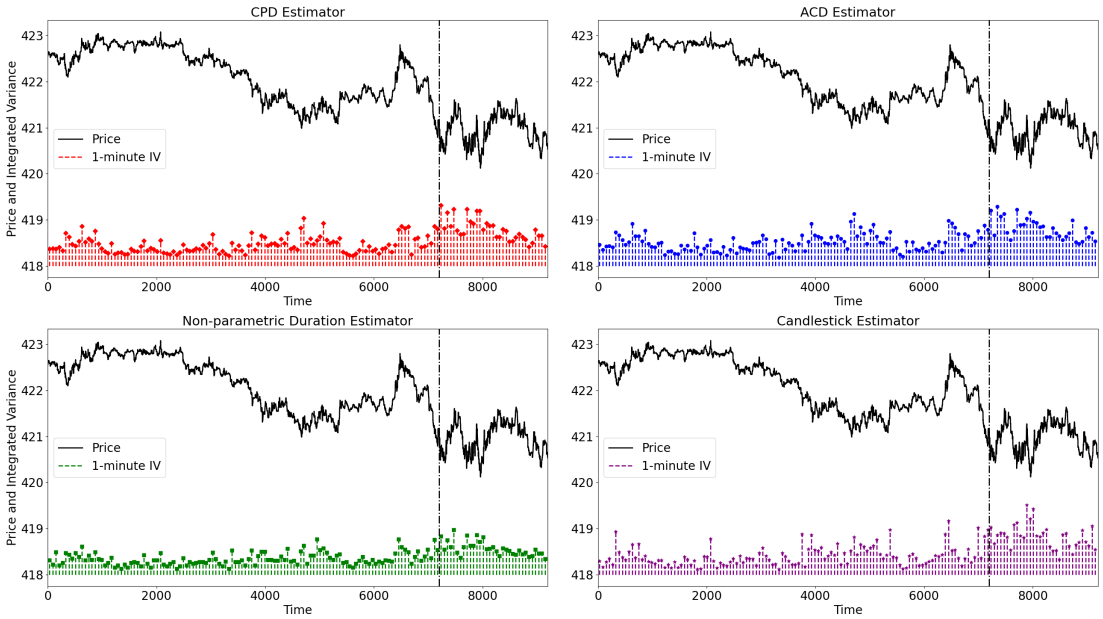

Figure 10. Estimated integrated variance over 1 min. Vertical dashed line marks the time of Boston Marathon bombing event.

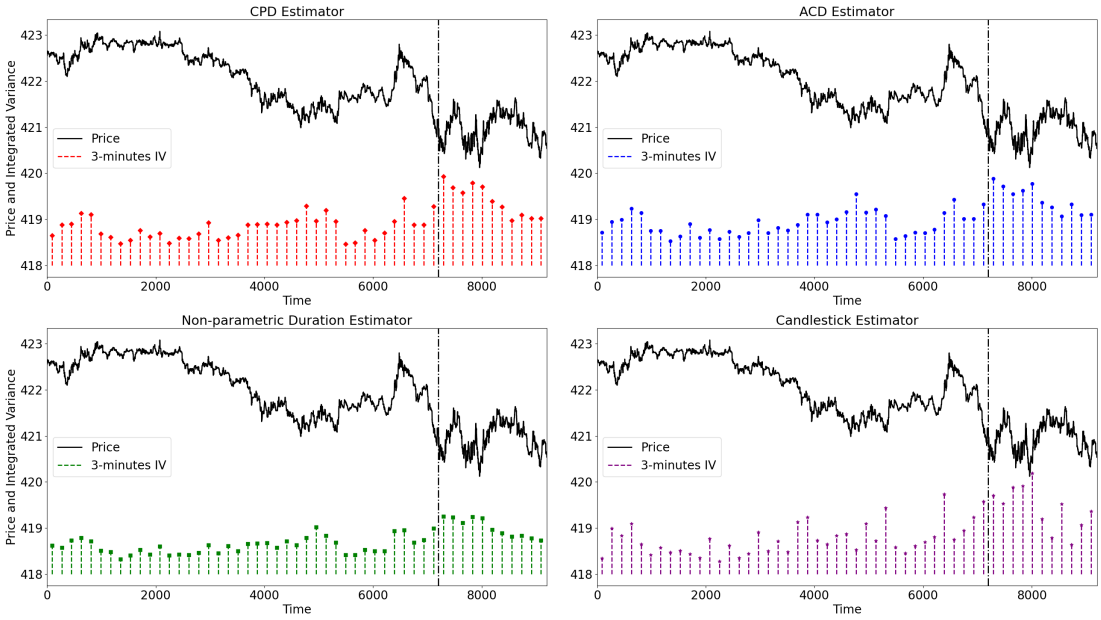

Figure 11. Estimated integrated variance over 3 min. Vertical dashed line marks the time of Boston Marathon bombing event.

5. Conclusions

This paper proposes the CPD model as a new framework for estimating spot volatility from price durations and demonstrates its advantages in high-frequency settings. Compared with alternative approaches such as the ACD model, the nonparametric duration estimator, and candlestick-based measures, the CPD model provides a more stable yet flexible characterization of volatility by accommodating structural breaks and sudden market shocks. Simulation results show that the CPD estimator closely tracks true volatility and captures regime shifts, while empirical analysis of Nasdaq limit order book data highlights its robustness in real markets. Particularly, it outperforms the benchmarks by capturing both persistence and responsiveness in high-frequency variance estimation in short intervals, and the application to the Boston Marathon bombing illustrates the model's ability to detect abrupt jumps in volatility while avoiding excessive noise. The model also offers practical value for high-frequency traders and market makers, for whom accurate and timely volatility estimation is vital to managing risk, setting spreads, and maintaining profitability. Overall, this work contributes to the market microstructure literature by introducing a novel duration-based approach that links order book dynamics to volatility estimation in ultra-high-frequency settings.

Nevertheless, this study has some limitations. First, the analysis is restricted to univariate volatility dynamics, whereas multivariate extensions would be necessary to model co-movements across assets. Second, the current CPD model can identify shifts in volatility but cannot predict the direction of jumps using additional order book information. Third, estimation of the CPD model is relatively computationally intensive, which constrains its applicability for real-time market monitoring. Addressing these limitations opens up promising avenues for future research. Extensions could include multivariate CPD models for systemic risk analysis, integration with liquidity and order-flow variables, and development of real-time estimation methods suitable for live trading systems. Exploring such directions would further enhance the applicability of the CPD framework for both academic research and practical risk management in high-frequency markets.

Author Contributions: Conceptualization, Z.L.; methodology, Z.L. and H.X.; software, Z.L. and Y.W.; validation, Z.L., H.X. and Y.W.; formal analysis, Z.L.; investigation, Y.W.; resources, Z.L. and Y.W.; data curation, Z.L. and Y.W.; writing—original draft preparation, Z.L.; writing—review and editing, Y.W.; supervision, H.X.; funding acquisition, Z.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Provincial Natural Science Foundation of Hunan Province, China, 2023JJ40191, which is acknowledged by Zhicheng Li.

Informed Consent Statement: Not appliable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Supplementary EM Steps

We refer to the our previous paper Z. Li and Xing (2022) to estimate the CPD model because the change-point structure is very similar to that work.

For the CPD model introduced in Section 2.2, the complete log-likelihood function is given by

$$\ell(\{y_1^n, \lambda_1^n\}) = \log P(\lambda_1) + \sum_{t=1}^{n} \log f(y_t \mid \lambda_t) + \sum_{t=2}^{n} \log P(\lambda_t \mid \lambda_{t-1}), \tag{A1}$$

where $\{y_t\}_{t=1}^n$ denotes the observed sequence of price-change durations and $\{\lambda_t\}_{t=1}^n$ the unobserved sequence of underlying intensities. Here, $P(\lambda_1)$ is the prior distribution of the initial intensity, $f(y_t \mid \lambda_t)$ the density of y_t conditional on λ_t, and $P(\lambda_t \mid \lambda_{t-1})$ the transition probability of λ_t given λ_{t-1}.

In our specification, the initial intensity λ_1 is independently drawn from a Gamma distribution $G(\cdot)$. Conditional transitions follow a Markov change-point process: with probability p, the intensity renews from $G(\cdot)$, and with probability $1 - p$, it remains unchanged. Thus, the log-likelihood simplifies to

$$\begin{aligned}\ell(\{y_1^n, \lambda_1^n\}) = \log G(\lambda_1) + \sum_{t=1}^{n} \log f(y_t \mid \lambda_t) \\ + \sum_{t=2}^{n} \Big[\log G(\lambda_t) \cdot \mathbf{1}_{(\mathbf{I_t=1})} + \log p \cdot \mathbf{1}_{(\mathbf{I_t=1})} + \log(1-p) \cdot \mathbf{1}_{(\mathbf{I_t=0})} \Big],\end{aligned} \tag{A2}$$

where $I_t = 1$ indicates that a change point occurs at the t-th trade, i.e., $\lambda_t \neq \lambda_{t-1}$. $\mathbf{1}_{(\mathbf{I_t=1})}$ is a indexing function that returns 1 when $I_t = 1$ and otherwise returns 0. Similarly, $\mathbf{1}_{(\mathbf{I_t=0})}$ equals 1 if there is no change-point at t-th event.

In the E-step, we compute the expected log-likelihood conditional on the observed durations $\{y_t\}$ and the current parameter estimates. In the M-step, we update (α, β, p) by maximizing this expected log-likelihood. The procedure iterates until convergence, yielding both parameter estimates and posterior distributions for the latent intensities.

Appendix A.1. Expected Likelihood

In the E-step, the expected log-likelihood, conditional on $\mathcal{D} = \{y_{1n}$, parameters of last iteration$\}$, is as follows:

$$\begin{aligned} \mathrm{E}\Big(l_c(\{y_{1n}, \lambda_{1n}\})\Big|\mathcal{D}\Big) &= \mathrm{E}\Big(\log G(\lambda_1)\Big|\mathcal{D}\Big) + \sum_{t=1}^{n} \mathrm{E}\Big(\log f(y_t|\lambda_t)\Big|\mathcal{D}\Big) \\ &+ \sum_{t=2}^{n} \mathrm{E}\Big(\log G(\lambda_t) \cdot \mathbf{1}_{(\mathbf{I_t}=\mathbf{1})}\Big|\mathcal{D}\Big) \\ &+ \sum_{t=2}^{n} [\log p \cdot P(I_t = 1|\mathcal{D}) + \log(1-p) \cdot P(I_t = 1|\mathcal{D})], \end{aligned} \tag{A3}$$

where the expectation is taken over the posterior distribution of hidden variables λ_i.

As shown in Appendix B, the posterior distributions of λ_t is as follows:

$$f(\lambda_t|\mathcal{D}) = \sum_{1 \le i \le t \le j \le n} \Pi_{itj} \cdot g_{ij}(\lambda_t), \tag{A4}$$

where $g_{ij}(\lambda) \sim \text{Gamma}(\alpha + (j-i+1),\ \beta + \sum_i^j y_t)$ and Π_{itj} are the change-point probability that the last change point occurs at i and the next change point occurs at $j+1$, which will be calculated in Appendix B.

Moreover,

$$P(I_{t+1} = 1|\mathcal{D}) = \sum_{1 \le i \le t} \Pi_{itt} \qquad P(I_{t+1} = 0|\mathcal{D}) = 1 - P(I_{t+1} = 1|D) \tag{A5}$$

where $t \in [1, N-1]$, and we also set $P(I_1 = 1|\mathcal{D}) \equiv 1$.

Therefore,

$$\begin{aligned} \mathrm{E}(\log f(y_t|\lambda_t)|\mathcal{D}) &= \int_{\lambda_t} \log f(y_t|\lambda_t) \cdot f(\lambda_t|\mathcal{D}) d\lambda_t \\ &= \int_{\lambda_t} (\log \lambda_t - \lambda_t y_t) \cdot \sum_{1 \le i \le t \le j \le n} \Pi_{itj} \cdot g_{ij}(\lambda_t) \cdot \mathrm{d}\lambda_t \end{aligned} \tag{A6}$$

$$\mathrm{E}(\log G(\lambda_t) \cdot \mathbf{1}_{(\mathbf{I_t}=\mathbf{1})}|\mathcal{D}) = \int \log G(\lambda_t) \cdot f(\lambda_t, I_t = 1|\mathcal{D}) \mathrm{d}\lambda_t$$

Since $f(\lambda_t|\mathcal{D}) = \sum_{1 \le i \le t \le j \le n} \Pi_{itj} \cdot g_{ij}(\lambda_t)$, thus $f(\lambda_t, I_t = 1|\mathcal{D}) = \sum_{t \le j \le n} \Pi_{ttj} \cdot g_{tj}(\lambda_t)$. Hence, we have the following.

$$\mathrm{E}(\log G(\lambda_t) \cdot \mathbf{1}_{(\mathbf{I_t}=\mathbf{1})}|\mathcal{D}) = \sum_{t \le j \le n} \Pi_{ttj} \int \log G(\lambda_t) \cdot g_{tj}(\lambda_t) \mathrm{d}\lambda_t \tag{A7}$$

Maximization and the Update of Parameters

Once we have the expected log-likelihood in (A3) and write it as a function form with the following arguments (α, β, p),

$$l_{EC}(\alpha, \beta, p) \equiv \mathrm{E}\Big(l_c(\{y_{1n}, m_{1n}\})\Big|\mathcal{D}\Big), \tag{A8}$$

we can perform our maximization step in EM and update our estimations of model parameters.

As only the last two items in (A3) contain parameter p, therefore, by first order maximization, we have the following.

$$\begin{aligned} \frac{\partial l_{EC}(\alpha, \beta, p)}{\partial \hat{p}} &= \frac{1}{\hat{p}} \sum_{t=2}^{n} P(I_t = 1|D) - \frac{1}{1-\hat{p}} \sum_{t=2}^{n} P(I_t = 0|D) \\ &= 0, \end{aligned}$$

which gives

$$\hat{p} = \frac{\sum_{t=2}^{n} P(I_t = 1|D)}{n-1}. \tag{A9}$$

Consider items in $l_{EC}(\alpha, \beta, p)$ contain parameters. It can be derived that (α, β)

$$\begin{aligned} & \mathrm{E}\Big(\log G(\lambda_1)\Big|\mathcal{D}\Big) + \sum_{t=2}^{n} \mathrm{E}\Big(\log G(\lambda_t) \cdot \mathbf{1}_{(\mathbf{I_t}=\mathbf{1})}\Big|\mathcal{D}\Big) \\ = \quad & [\alpha \log \beta - \log \Gamma(\alpha)] \cdot A + (\alpha - 1) \cdot B - \beta \cdot C, \end{aligned} \tag{A10}$$

where

$$A = \sum_{t=1}^{n} \left(\sum_{t \le j \le n} \Pi_{ttj} \right) = \sum_{t=1}^{n} P(I_t = 1|\mathcal{D})$$

$$B = \sum_{t=1}^{n} \left[\sum_{t \le j \le n} \Pi_{ttj} \cdot \int_{\lambda_t} \log \lambda_t \cdot g_{tj}(\lambda_t) \mathrm{d}\lambda_t \right]$$

$$C = \sum_{t=1}^{n} \left[\sum_{t \le j \le n} \Pi_{ttj} \cdot \int_{\lambda_t} \lambda_t \cdot g_{tj}(\lambda_t) \mathrm{d}\lambda_t \right].$$

Therefore, by maximizing (A10), we have

$$\hat{\alpha} = \frac{C}{A} \hat{\beta}, \tag{A11}$$

and

$$A \cdot \log \hat{\beta} - A \cdot \psi(\hat{\alpha}) + B = 0. \tag{A12}$$

Thus, combining (A11) and (A12), we can solve the new value of $\hat{\alpha}$ and $\hat{\beta}$. In addition to the result of $\hat{p}$ in (A9), we can perform the next EM iteration until the convergence of estimators.

Appendix B. The Posterior Distribution of λ_t in EM

Appendix B.1. Forward–Backward Filter

This section outlines the forward–backward filtering procedure for updating the posterior distribution of λ_t given $\{y_1, \ldots, y_n\}$.

Forward Filter

Let $R_t = \max\{k \mid \mathbf{I}_k = 1, k \le t\}$ denote the most recent change point before or at t. Given $(y_{1t}, R_t = s)$, the density of the posterior distribution of λ_t is

$$g_{st}(\lambda_t) \triangleq f(\lambda_t | y_t, R_t = s) \propto \prod_{i=s}^{t} f(y_i|\lambda_t) G(\lambda_t). \tag{A13}$$

The posterior distribution of λ_t given y_{1t} can be expressed as

$$f(\lambda_t | y_{1t}) = \sum_{i=1}^{t} p_{it} \cdot g_{it}(\lambda_t), \tag{A14}$$

where $p_{it} = P(R_t = i|y_{1t})$ are mixture weights calculated recursively.

Backward Filter

Define $\widetilde{R}_{t+1} = \min\{k \mid I_{k+1} = 1, k \ge t+1\}$, i.e., the next change point after t. Let $q_{t+1,j} = P(\widetilde{R}_{t+1} = j|y_{t+1,n})$.

Then, the posterior distribution of λ_t given $y_{t+1,n}$ is

$$f(\lambda_t|y_{t+1,n}) = p \cdot G(\lambda_t) + (1-p) \sum_{j=t+1}^{n} q_{t+1,j} \cdot g_{t+1,j}(\lambda_t), \tag{A15}$$

where the weights $q_{t+1,j}$ can be obtained recursively.

Combination (Forward–Backward Algorithm)

The posterior distribution of λ_t given the full sample $\{y_1, \ldots, y_n\}$ is

$$f(\lambda_t|y_{1n}) = \sum_{1 \le i \le t \le j \le n} \Pi_{itj} \cdot g_{ij}(\lambda_t), \tag{A16}$$

where Π_{itj} is the joint posterior probability of the last change point at i and the next at $j+1$.

Appendix B.2. Calculation Steps of the Posterior Intensity Distribution

In practice, the posterior intensity distribution can be evaluated as follows:

- The conditional densities $g_{ij}(\lambda)$ are Gamma distributions, i.e., $g_{ij}(\lambda) \sim \text{Gamma}(\alpha + (j-i+1),\ \beta + \sum_{t=i}^{j} y_t)$.
- The next thing is to calculate Π_{itj}. By the foward–backward filter part, we can derive that $\Pi_{itj} = \dfrac{\Pi^*_{itj}}{\sum\limits_{1 \le s \le t \le k \le n} \Pi^*_{stk}}$, and

$$\Pi^*_{itj} = \begin{cases} p \cdot p_{it} & j = t,\ 1 \le i \le t \\ (1-p) p_{it} \cdot q_{t+1,j} \cdot \frac{f_{ij}}{f_{it} f_{t+1,j}} & j > t,\ 1 \le i \le t, \end{cases} \tag{A17}$$

  where f_{ij} is defined as

$$f_{ij} = \int \prod_{t=i}^{j} f(y_t|\lambda) \cdot G(\lambda) d\lambda.$$

- The normalizing constants f_{ij} are obtained in closed form due to conjugacy:

$$f_{ij} = \frac{\Gamma(\alpha + j - i + 1) \cdot \beta^{\alpha}}{\Gamma(\alpha) \cdot (\beta + \sum_{t=i}^{j} y_t)^{\alpha + j - i + 1}}.$$

- The forward probabilities p_{it} and backward probabilities $q_{t+1,j}$ are computed recursively, which in turn yield Π_{itj} and hence the posterior $f(\lambda_t|\mathcal{D})$.

In summary, the forward–backward algorithm shows that the posterior of λ_t is a finite mixture of Gamma distributions, with weights updated recursively. This structure makes the EM algorithm computationally tractable for high-frequency data.

Appendix C. Comparison with TACD Model

Here, we compare our model with the nonlinear threshold ACD (TACD) model proposed by Zhang et al. (2001), using the Apple LOB data from 15 April 2013. The results are plotted in the figures below. From the comparison, we can still find that (1) the residuals from the CPD model conform closely to the exponential distribution, while those from the TACD model show substantial deviations; (2) the CPD model produces volatility estimates that are more persistent and less noisy, whereas the TACD model generates highly volatile series with frequent sharp spikes; (3) the CPD-based integrated variance estimator captures structural shifts in volatility more effectively, while the TACD estimator tends to exaggerate short-term fluctuations and overreact to abrupt price changes.

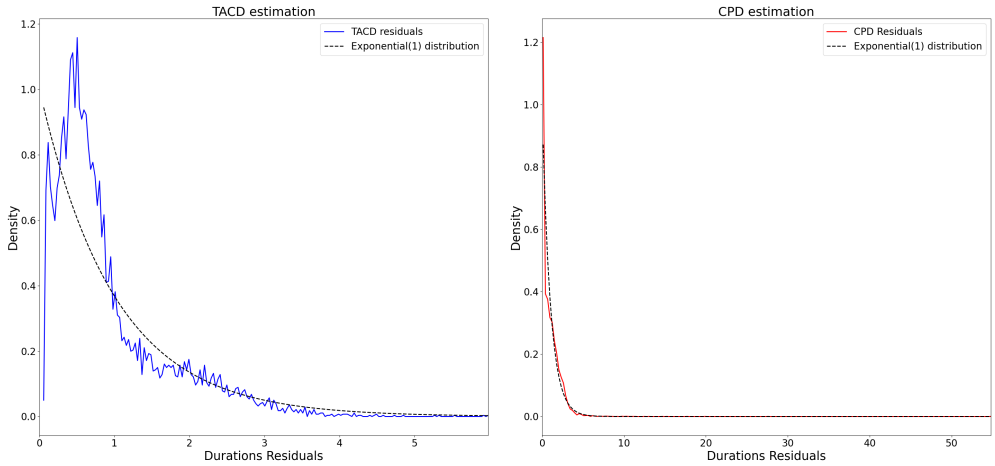

Figure A1. Fitted duration residuals by TACD model and CPD model.

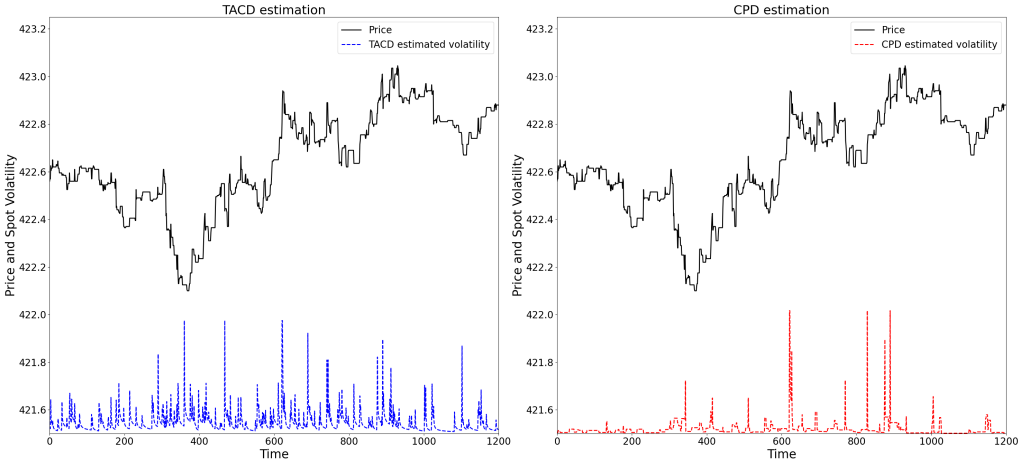

Figure A2. Spot volatility estimation by TACD model and CPD model.

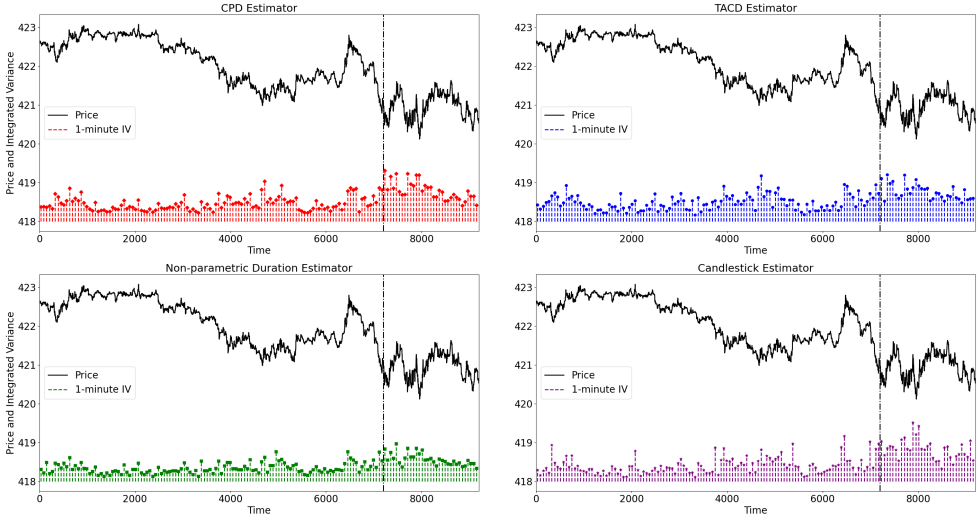

Figure A3. Estimated integrated variance over 1 min. Vertical dashed line marks the time of Boston Marathon bombing event.

Notes

1 These approaches build on earlier work by Gallant et al. (1999) and Alizadeh et al. (2002), who introduced range-based volatility estimators for modeling and forecasting time-varying volatility.

2 Boston Marathon bombing occured on 15 April 2013. https://en.wikipedia.org/wiki/Boston_Marathon_bombing (accessed on 8 July 2025).

References

Admati, A. R., & Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies*, *1*(1), 3–40. [CrossRef]

Alizadeh, S., Brandt, M. W., & Diebold, F. X. (2002). Range-based estimation of stochastic volatility models. *The Journal of Finance*, *57*(3), 1047–1091.

Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, *39*(4), 885–905. [CrossRef]

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, *71*(2), 579–625. [CrossRef]

Barndorff-Nielsen, O. E., & 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. [CrossRef]

Barndorff-Nielsen, O. E., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, *2*(1), 1–37. [CrossRef]

Barndorff-Nielsen, O. E., & Shiryaev, A. N. (2015). *Change of time and change of measure*. World Scientific Publishing Company.

Bauwens, L., & Giot, P. (2000). The logarithmic ACD model: An application to the bid-ask quote process of three NYSE stocks. *Annales d'Economie et de Statistique*, *60*, 117–149.

Bollerslev, T., Li, J., & Li, Q. (2024). Optimal nonparametric range-based volatility estimation. *Journal of Econometrics*, *238*(1), 105548. [CrossRef]

Box, G. E., & Tiao, G. C. (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical association*, *70*(349), 70–79.

Cho, D. C., & Frees, E. W. (1988). Estimating the volatility of discrete stock prices. *The Journal of Finance*, *43*(2), 451–466. [CrossRef]

Daley, D. J., & Vere-Jones, D. (2003). *An introduction to the theory of point processes: Volume I: Elementary theory and methods*. Springer.

Easley, D., Engle, R. F., O'Hara, M., & Wu, L. (2008). Time-varying arrival rates of informed and uninformed trades. *Journal of Financial Econometrics*, *6*(2), 171–207. [CrossRef]

Easley, D., Hvidkjaer, S., & O'hara, M. (2002). Is information risk a determinant of asset returns? *The Journal of Finance*, *57*(5), 2185–2221. [CrossRef]

Easley, D., & O'hara, M. (1992). Time and the process of security price adjustment. *The Journal of Finance*, *47*(2), 577–605. [CrossRef]

Engle, R. F., & Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, *66*(5), 1127–1162. [CrossRef]

Gallant, A. R., Hsu, C.-T., & Tauchen, G. (1999). Using daily range data to calibrate volatility diffusions and extract the forward integrated variance. *Review of Economics and Statistics*, *81*(4), 617–631. [CrossRef]

Gerhard, F., & Hautsch, N. (2002). Volatility estimation on the basis of price intensities. *Journal of Empirical Finance*, *9*(1), 57–89. [CrossRef]

Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, *14*(1), 71–100. [CrossRef]

Hautsch, N. (2011). *Econometrics of financial high-frequency data*. Springer Science & Business Media.

Hong, S. Y., Nolte, I., Taylor, S. J., & Zhao, X. (2023). Volatility estimation and forecasts based on price durations. *Journal of Financial Econometrics*, *21*(1), 106–144. [CrossRef]

Hujer, R., Vuletic, S., & Kokot, S. (2002). *The markov switching acd model* (Tech. Rep.). Working paper series: Finance & Accounting. Johann Wolfgang Goethe-Universität Frankfurt am Main, Fachbereich Wirtschaftswissenschaften

Lai, T. L., Liu, H., & Xing, H. (2005). Autoregressive models with piecewise constant volatility and regression parameters. *Statistica Sinica*, *15*(2), 279–301.

Li, J., Wang, D., & Zhang, Q. (2024). Reading the candlesticks: An OK estimator for volatility. *Review of Economics and Statistics*, *106*(4), 1114–1128. [CrossRef]

Li, Z., & Xing, H. (2022). High-frequency quote volatility measurement using a change-point intensity model. *Mathematics*, *10*(4), 634. [CrossRef]

Tse, Y.-K., & Yang, T. T. (2012). Estimation of high-frequency volatility: An autoregressive conditional duration approach. *Journal of Business & Economic Statistics*, *30*(4), 533–545. [CrossRef]

Zhang, M. Y., Russell, J. R., & Tsay, R. S. (2001). A nonlinear autoregressive conditional duration model with applications to financial transaction data. *Journal of Econometrics*, *104*(1), 179–207. [CrossRef]